



THREAT

MODELING

FUNDAMENTALS



HAKAN
GEIJETZ



War is the realm of uncertainty; three quarters of the factors on which action in war is based are wrapped in a fog of greater or lesser uncertainty. A sensitive and discriminating judgment is called for; a skilled intelligence to scent out the truth.

Carl von Clausewitz, *On War*, 1832

INTRODUCTION

There is tension anarchists feel between the steps we believe we should take against all forms of domination and the consequences we are willing to risk in pursuit of those goals. Repression, either explicitly from the State or implicitly from non-State actors, constrains the set of actions we are willing to take, and it primarily does so by setting a low ceiling on the “extremeness” of those actions. Extreme actions drive change, and agents of the State know this, so they do what they can to take away our most powerful tools.

But it is possible to reclaim them.

We turn to operational security (OpSec) and security culture as the primary means of “getting away with it,” or more precisely put: reaching our goals with minimal consequences. These terms are used quite casually, and the practices around them can at times be both opaque and dogmatic even when they contribute to reducing the effects of repression. The phrase “threat modeling” gets thrown around with even less explanation of what it means or how one goes about doing it.¹ What literature there is on threat modeling tends to be overwhelmingly focused on securing corporate IT systems against hackers, and while there are interesting lessons one can learn from these texts, they require a significant lateral move to be able to apply them to the average radical on the streets. This zine aims to fill that gap.

Threat modeling provides the justification for the various practices of OpSec and security culture (henceforth just “security” for brevity). Someone might tell you to leave your mobile phone at home for an action, and this isn’t just for fannies, but because phones—even when powered off—leak location data. Every security practice and norm should have an evidence-informed threat model behind it with traceability from the observed and inferred actions of the adversaries to the countermeasures taken against them. Some practices that have become outdated remain because of tradition, and new practices that should be adopted often aren’t because people don’t understand the threat landscape in which they operate. Threat modeling is how these flaws are identified and resolved.

Threat modeling can sound like a niche expert field, but it’s something all of us do every day. You might find crafty ways to slack off at work and do so by trying

¹As for what a “threat” even is, for now just think of something an enemy might intentionally do to harm you, but we’ll see a slightly more precise definition later.

to maximize the amount of faffing about you can do before getting caught. Your boss's disposition, the presence of security cameras, and tattletale coworkers might influence your actions, and the amount you slack off might change over time or even depending on which shift you're working. Every day, we ask ourselves what might happen, how likely it is, and what we can do about it, and we adjust our behavior.

This zine is written to be accessible by everyone, not just those who already have an interest in security. It assumes that you have no knowledge of security in the context of radical social movements, but also seasoned veterans will find use in a more structured discussion of the security practices they already apply. While the general principles of threat modeling that are discussed here can be applied to many scenarios, this zine focuses specifically on resisting repression from local law enforcement, intelligence agencies, and fascists both organized and lone wolves.

This zine is not a singular authority on how to threat model or what one's security protocol should look like. Every person is different, every scene has its quirks, and every region has its unique threats of cops, fascists, and other dastardly villains. These all change with time. Take what you can from this text, adapt it, and leave the useless or outdated parts behind.

A small amount of threat modeling can do a great deal to decrease the effects of repression while increasing the range of possible strategies one can use in pursuit of their goals. It can be done as a solo exercise, during casual conversation with comrades, or as part of a focused analysis in prep for a major action. Once you have learned this skill, you can be more confident in your ability to reduce the effects of repression and you can decrease the strain and overhead when planning actions or even just existing as a radical.

VOCABULARY

Threat modeling is a structured process. Thus, we first need a well-defined vocabulary to ensure a shared understanding of the words we're using.

Subject — A subject is a person or group who may be the target of scrutiny, repression, or espionage. It is the entity the threat model concerns. You might be the subject of your own threat model as well as (directly or indirectly) the subject of one of your crew's threat models.

Goal — A goal is something the subject wants to achieve.² Goals might be things like “disrupt nazi group X” or “avoid getting doxxed.”

Strategy — A strategy is a specific set of actions taken by a subject to achieve their goals. It is the exact path among many possible paths to reach one (or several)

²Many texts about threat modeling talk about assets that a subject wants to protect, and this makes sense when discussing valuables in a safe or data in computer networks. For radicals, we're generally less concerned with physical assets but with intangible ones. Hence, we use the term “goals.”

of their many possible goals. If you do classic antifascism and your goal is to disrupt fash organizing, a strategy you might select is: next Tuesday night, we are going to place 500 flyers outing a local chud throughout the neighborhoods where they both live and work.

Adversary — An adversary is a person or group that wants to prevent a subject from realizing their goals. Adversaries can be internal (snitches, grifters) or external (cops, right-wing street goons). They can be direct (cops) or indirect (false allies, competing factions).

Capability — A capability is knowledge, a skill, or an item an adversary has that they may use against a subject to prevent them from achieving their goals. Capabilities might be a fleet of motorbikes, the monitoring of live internet traffic from an ISP, or the ability to use legal or extralegal violence.

Vulnerability — A vulnerability is an aspect of a subject's life or security protocol that can be exploited by an adversary to disrupt their goals or retaliate against them.³ A habit of bragging is a vulnerability as it might cause a subject to leak information about past secret actions. Lacking citizenship in the State in which the subject resides (i.e., being able to be deported) can also be a vulnerability.

Threat — A threat is realistic chance that an adversary exploits a vulnerability. It can be abstract like “someone might hack your computer” or concrete like “nazi group X will turn up at the next drag event at venue Y.” The adjective “realistic” is included to keep our scope of investigation narrow. While it is within the State's capabilities to drone strike you on your way to work, if you are living in the so-called West, the odds of this—at this time—are near zero, so it is not actually a threat.

Impact — Impact is a measure of the negative consequences if a vulnerability is exploited. A vulnerability that could reveal a subject's biometric information during an action may have high impact (decades imprisonment). A vulnerability that could reveal a subject's typical working hours may have low impact (most people work during the day/evening).

Probability (of Impact) — The probability of impact is a measure that considers both how likely it is that an adversary will attempt a threat and how likely it is that it will succeed.⁴ Surveillance cameras are near certain to be present on interesting properties, and without countermeasures against the such as covering your face or tattoos, they have high probability of having impact. A rent-a-cop happening to drive down a street right as you tag a building might, due to their sheer infrequency, have low probability of impact even if them spotting you is near certainty of getting caught.

³Other literature differentiates between weaknesses and vulnerabilities, and in the context of IT systems this might make sense. A program can have a design weakness that is not a vulnerability. It makes no sense to claim a human has a design weakness in how they live their life under repression.

⁴Some methods for threat modeling will use two measures (exploitability and likelihood [of attempted threat]), but for simplicity in this zine, they are reduced to one measure. If you want to split these back up for use in ranking threats, by all means do so.

Risk — Risk is the combined measure of impact and it’s probability. A vulnerability with extremely high impact that has a realistic but tiny probability of occurring may be considered low risk, but a vulnerability with only moderate impact but high probability might be considered high risk.

Countermeasure — A countermeasure is an action taken to reduce risk.⁵ Countermeasures may work to reduce probability by addressing the vulnerability itself, or they may address impact by altering adjacent areas of the subject’s life or security protocol. Countermeasures do not need to make risk zero to be worthwhile.

Security Protocol — A security protocol is the set of countermeasures taken by a subject given a defined set of adversaries with certain known or assumed capabilities.⁶ It might mean the specific OpSec strategies taken during an action and its preparatory phase, or it might mean the norms that constitute the security culture of a given milieu.

Threat Model — A threat model is the output of the process of threat modeling. It is a model that enumerates a subject’s goals and strategies, their adversaries and capabilities, and the countermeasures the subject can use against them. The threat model informs a security protocol that guides the subject’s actions. Sometimes the term is used to mean a specific adversary with known capabilities, for example “our threat model is [against] State domestic intelligence services.”

THE BASICS

Threat Modeling is the structured process of identifying threats to your goals and selecting countermeasures that can be deployed against these threats. The goal of threat modeling is to analyze your behaviors and strategies to learn how you have or might expose yourself to repression. There are many ways to threat model, and they all have their advantages and disadvantages. There are some characteristics that all tend to share. Most threat modeling methods in some way answer the following questions:⁷

1. What are our goals?
2. How might someone oppose us in them?
3. What can we do about it?
4. Is our threat model predictive? And is our security protocol effective?

The last point is key because threat modeling isn’t just iterative during a single

⁵In other literature, these are called “mitigations,” but due to the extremely active nature of OpSec versus simply updating a computer, the term countermeasures seems more appropriate.

⁶This is sometimes called “a security plan,” but plans are flexible and often disregarded on a whim. A protocol is something both rigid and demanding, and these are positives. If it is too rigid, it needs to be collectively renegotiated, not ignored at random without alerting others.

⁷This is adapted from the Four Question Framework.

session. It is also iterative over time as its successes or failures become apparent when it is applied to the real world.

When threat modeling is done well, there are two desirable outcomes: new threats or countermeasures are discovered, and the model trends towards better predictive power.⁸ Even for the most experienced crews, threat modeling should reveal something new. If it hasn't (i.e., if you're only writing down what you already know), the exercise can still be beneficial in as much as it ensures a shared threat model and security protocol. Even so, if there are no novel discoveries from threat modeling, more research should be done to uncover unknown threats or superior strategies.

The end product of iterated threat modeling needs to be fairly specific. "Create anarchy" is not a specific enough goal to be actionable, nor is "someone might spy on us" a specific enough adversary and capability to defend against. Specificity is important not only because it informs us about what we need countermeasures against, but also what we *don't* need countermeasures against. Non-specific goals or adversaries can be starting points, but you will need to iterate until they becomes specific.

You can threat model alone. Doing so can be helpful because it will allow you to identify the limits to your tolerance for risk as well as the sorts of goals you actually want to pursue. Having goals that do not align with those of your comrades means someone might be compromising their ideals or desires in order to work with others. Differing tolerances for risk can mean someone will feel anxious—possibly to the point of becoming unreliable—or that someone will feel they aren't doing enough and should be taking bolder action. It is often easier to meditate on your goals and risk tolerance on your own than in front of a group where social pressures and bravado can influence you to hide your true preferences. Once you have your own model, go to your crew and threat model with them. You may learn that you need to find a new affinity group, and that's okay.

Threat modeling in a group has advantages of being able to use others' knowledge to inform the model. Most of what we do is with at least one other person whether it's handing out pamphlets or making drilling equipment unusable. Individual threat models can have a limit to how well they can inform us because our goals might not actually be the goals of our crew. The security protocol that comes out of collective threat modeling might be something that you aren't comfortable with. You might agree with the crew's goals and strategies, but might find their slipshod approach to security creates an intolerable amount of risk for you.

You need to prepare for threat modeling. You probably want to dedicate an hour to this if you're doing it alone and two to four hours if you're doing it with a crew. You may need to do multiple sessions because of research that's needed to

⁸Note that this doesn't say "accurate." All models have some intentionally built-in inaccuracies because infinite precision would lead to models complex beyond human understanding. Accuracy to the extend possible is still important as an inaccurate model that wholly does not reflect reality will be unable to make meaningful predictions about the actions and reactions of one's adversaries.

fill knowledge gaps.

There is a circular dependency on threat modeling. In order to know *where* and *how* it's even safe to threat model, one has to first have a threat model that answers the question: would the State repress us for even having such a discussion?⁹ There may already be the limitation that places frequented by radicals (your home, a squat, an infoshop, etc.) might already under surveillance and could be unsuitable for such discussions. You may need to do it without electronic devices at an unfamiliar cafe or park. Bring paper and something to write with. You may also need to bring a lighter because depending on what you're threat modeling, you may need to destroy the model itself and only save the resulting security protocol.

ONE PARTICULAR METHOD

The method described in this section is not “the best,” even to the extent such a thing might even exist. This is just one method that will be discussed in depth enough to allow you to design your own. I'm not even attempting to name it to avoid giving it excess importance relative to whatever methods you might come up with.

This method is goal-oriented. It focuses on what you want to do then applies constraints.¹⁰

IDENTIFY GOALS

If your goals or desires are unclear, the first step is to materialize them into something specific. Start by brainstorming or making a mind map. Write down every idea you have on a piece of paper, or alternatively write down each idea on a separate note card. Don't worry about feasibility or what can go wrong. Just start writing. Take goals and break them down in to sub-goals or write down prerequisites you need to reach that goal. Draw connections between related goals or sub-goals, or if you're using note cards, cluster them with each other.

For some individuals, their rough goals and methods are already honed from having been part of radical movements for many years. An affinity group might be operating under the guiding principle of disrupting fascists in their immediate vicinity. Time and material constraints might limit their work, so their goals are scoped to what they can reasonably achieve while still managing to feed themselves and find some joy in life. Identifying goals may be more akin to target selection.


⁹This applies generally for all adversaries. If you're threat modeling leaving an abusive relationship, your partner might spy on you to prevent you from attempting to leave.

¹⁰This is opposed to, say, adversary-oriented threat modeling that focuses on adversaries and their capabilities then looks at what goals are left available given those constraints. A weakness of the adversary-oriented approach is that a few infrequent cases of shows of force or lucky breaks in investigation can cause us to take those as a baseline of their capabilities. This preemptively cuts off certain avenues of attack that we may actually be able to leave open. As we'll see in the coming sections, assessing the probability of a threat needs to come into play, not just the possibility that *it could happen at all*.

Goals don't have to be "classic" anarchist goals like attacking this or building that. Depending on one's life circumstances or identities, a goal might be simply to survive and thrive in spite of oppression. This might lead to sub-goals of avoiding interactions with police, not drawing attention to one's self, or gathering enough resources to move somewhere safer.

Because anarchism is a social movement and not a purely individual pursuit, we implicitly pull in goals that are altruistic towards others. Our strategies and security protocol should generally aim to protect others from legal entanglements, incarceration, and bodily or psychological harm. These altruistic goals exist so that our security protocol doesn't simply become "how do *I* avoid arrest" but rather "how do we *all* avoid arrest."

Figure 1: Flow Chart for Feeding the Homeless



IDENTIFY GOAL PREREQUISITES

Once you have your goals, the prerequisites need to be clearly captured. By writing down the steps you need to achieve that goal, you can investigate the goal's feasibility, and seeing the steps written out will let you find vulnerabilities in later stages of threat modeling. In most cases, there is not a single path to your goal, so write down all possible ways you can accomplish it. Basic mind maps and flowcharts work fine for this.¹¹

¹¹Some people use fishbone (Ishikawa) diagrams, but I find those hard to read, and they can get cluttered.

In the (simplified) example in Figure 1, there are two requirements for feeding a homeless person: we have to both have food and be able to find them. For both of these sub-goals, there are several ways to accomplish them.

IDENTIFY ADVERSARIES AND CAPABILITIES

Identifying goals can take some soul searching as we determine what matters to us or what strategies we consider to be effective or ethical. Identifying adversaries and capabilities takes genuine research.

“Common sense” tells us that as members of an anti-authoritarian social movement our general adversaries are local cops, national security agencies, individual right-wing lone wolves, and organized groups of fascists. This same common sense suggests the general shape of their capabilities such as collecting forensic evidence, wire tapping, or just plain ol’ violence. The things one might call “common sense,” another might call delusional paranoia, and yet another might dismissively call a simpleton’s idea of the State’s and fascists’ methods of operating. Teasing out what of this common sense reflects reality is why we need to do research on each of the claims.

While you may be tempted to write down every alphabet agency that exists in your region, there might be little meaningful difference between two national intelligence agencies or two local police detachments. There are only a handful of law enforcement units, and only a handful of fascist groups. They can typically be grouped into buckets like so:

- State: domestic intelligence, local police.
- Non-State: groups who strike first, groups who strike back, groups who are effectively non-violent.

However, with the use of data sharing and fusion centers, or more generally with the digitization of policing and use of computers to do inhuman amounts of data collating, capabilities might start to blur between intelligence agencies and local police, so these buckets might in the coming years cease to be as clear or even meaningful.¹²

Prune away buckets of adversaries where it is reasonable. As a soup kitchen, you probably aren’t under active investigation by domestic intelligence or the most likely target for deep-cover infiltration.¹³ Additionally, prune non-State entities. Rival factions might not attack each other, and for carrying out actions their existence may be totally irrelevant. Gangs or mafias might leave you well enough alone

¹²For example, the NYPD has the Department Intelligence Bureau which operates outside classical oversight and has foreign intelligence assets and direct connections to foreign police. Does this make them local police in the classical sense? An intelligence agency? Or something else entirely?

¹³They might use a soup kitchen to get a foothold to hop elsewhere in the scene, but the soup kitchen itself likely isn’t the primary target.

if you don't infringe on their racket, and if that's not your goal, they may not need to be a concern for your threat model.

USE A THREAT LIBRARY

Buckets of adversaries are identified so that specific capabilities can be tied to them. Two crews on opposite sides of the same State that are up to equally attention-drawing activities will have the same adversary with the same capabilities. Two crews in entirely different parts of the worlds may just as well be up against nearly identical adversaries (e.g., two non-allied States with similar domestic intelligence agencies), and as a result if these two crews enumerate their adversaries' capabilities, they will have very similar lists. This is to say that while threat modeling must be done individually, research about adversaries can be shared to massively reduce the amount of duplicated effort. Knowledge pools of this sort are called threat libraries.

A threat library categorizes and explains capabilities adversaries might have and the threats they produce. This could be a stack of color coded note cards you keep on hand so you can use them offline, or it could be a database hosted online somewhere.¹⁴ An example can be found in Table 1.

Table 1: Sample Index of a Threat Library

Physical Surveillance	PsyOps	Hacking
Landline wire taps	Spreading rumors	Hardware exploits
Surveillance cameras (outdoor, indoor)	Infiltrators (second order effects)	Backdoor deals with manufactures
Mobile phone location tracking	In-person harassment	Zero-days in open source libraries

Whether you use an external one or you create your own, there's a few features that threat library should have to be useful.

¹⁴The Counter Surveillance Research Center's threat library (csrc.link/threat-library) is the only one I've found that looks like it is tailored towards radicals and is of passing quality. My general assessment of it at this time is that it is certainly a good place to start one's research, but it is both insufficiently broad and insufficiently deep to be *completely* reliable for most use cases. They are open to outside collaboration as per their website, so it is my genuine hope that others contribute and make it into a high-quality resource.

- Threats should have clear and specific names.
- Threats should link to related threats including adjacent, parent, and child threats.
- Threat that have occurred in the past even if they no longer are known to be occurring should be included. They may occur again.
- Theoretical threats that are on the horizon even if they are known to not yet exist should still be included (especially digital ones).
- Threats should list clear examples of them having happened. This makes them “real” and also gives branch points for further research into case studies.
- Threats should be mapped to adversaries with some sort of notion of how frequently they are used or lead to an adversary reaching their goals (e.g., how often do national police who hack their targets’ computers get convictions based off this?).


CREATE ATTACK TREES

Like how your goals have sub-goals or prerequisites, your adversaries have the same. Cops’ goals are generally “stop anarchists from doing stuff” and “arrest anarchists after they do stuff.” A local group of chuds might have a simple goal of “beat up on queers.” Those are vague, and we might refine them into specific threats. These threats themselves can be refined to show how they might actually occur. A flow chart that shows prerequisites for a given threat is called an attack tree. An example can be found in Figure 2.

When creating attack trees, you may annotate each node with information to help you determine if a specific threat is plausible. One annotation is an estimate of expended effort for each prerequisite or sub-goal. If there are two ways to reach a sub-goal, your adversary will most likely use the one that take less effort. You could also annotate it with something like blowback. An extrajudicial—or quasi-legal due to police immunity—killing might have some blowback against the individual carrying out the act or the agency that ordered it, and you might estimate that avoiding blowback is something local cops do. If your adversary is non-State, you might annotate the sub-goals with their legality, or rather odds of prosecution due to cops letting chuds do their dirty work.

Attack trees can help you think like your adversary to help you guess what their strategies for disrupting you might be. They come with the limitation that you don’t actually know how your adversary thinks or how they appraise certain situations. You as an anarchist might assume some level of repression for engaging in criminalized acts that some fascistic actor does not. Estimating effort of complexity might also be difficult because you might overestimate an adversary’s technical competence or underestimate time/funds allocated to pursuing you. Like with the rest of your threat model, there might be things you completely fail to see.

Figure 2: Direct Action Investigation Attack Tree (Incomplete)



A related concept to the attack tree is that of the kill chain. The idea of the kill chain is that there is a sequence of events that must occur for your adversary to stop you from reaching your goal. If you can interrupt this sequence at any point (i.e., breaking the chain) your countermeasures will be sufficient to let you reach your goal. Because there might be a number of sequences, we can visualize them as tree. If the tree is accurate, then we can find ways to cut enough branches and boughs to prevent the adversary's goal from being reached.

ENUMERATE AND PRIORITIZE THREATS

Once you have a list of adversaries and a method for listing what threats they might pose, you need to somehow make this list actionable. If your threat library is sufficiently detailed and you've done enough research into similar crews, you will have a sense of what capabilities are actually deployed. However, you should keep in mind that your adversary may not actually be using all their capabilities in order to keep their true maximum capabilities unknown to the public. In the case of the State, police may use parallel construction so that even if you carefully read through all the legal cases and evidence, you might only know how they *said* they got enough evidence for a conviction, but you might not know how they *actually* got it.

After removing threats that have a negligible chance of happening, your list will probably still be too large for you to fully address. To prioritize threats, you need a heuristic. One method is to use "gut instinct" and just arrange them into what feels right. This is entirely reasonable. Another method is to assign each threat an

impact and a probability, then to multiply those scores together.

Scores for probability and impact can be found in Table 2. The numbers are skewed to place greater emphasis on impact and especially severe impacts. Be aware that “catastrophic” is relative to the goals being considered. Catastrophic for tagging (a moderate fine) versus catastrophic for sabotage (a decade in prison) might be wildly different and aren’t comparable in this context.

If after sorting the threats, you feel the ordering is “wrong,” you can rearrange them. The scores are just to help you do an initial sorting.

Table 2: Calculating Risk Scores

		Impact				
		None (0)	Minor (1)	Moderate (3)	Major (5)	Catastrophic (10)
Probabilities	Never (0)	0	0	0	0	0
	Rare (1)	0	1	3	5	10
	Unlikely (2)	0	2	6	10	20
	Likely (3)	0	3	9	15	30
	Almost Certain (5)	0	5	15	25	50

$$\text{Risk} = \text{Probability} \times \text{Impact}$$

Even the best of us, those with years of experience as radicals are limited in how accurately we can assess risk, or more specifically probability (of impact). For certain classes of actions, all probabilities might stack up around “rare” which means we’d change the names of probabilities to something like rare, less rare, rather uncommon, etc. How we might decide what is or isn’t what exact level of rare is, at best, based on not-particularly-accurate gut feelings. Estimating risk is no easy task.

IDENTIFY COUNTERMEASURES

Once threats are identified, they need to be resolved. To resolve them doesn’t mean to solve them in the sense of the problem has been completely handled. A resolution is just a conscious decision about what, if anything, is to be done about the particular threat. Threats can be considered resolved in one of four ways.

Accepted — A threat is accepted if it is decided that nothing can be done about it. The threat is identified, countermeasures are considered, and then if it is determined that the goal or strategy to which the threat applies is too critical to alter or remove or that the countermeasures are not feasible given current resources, the threat is marked as accepted.

Avoided — A threat is avoided if a goal or strategy is completely removed from the possible courses of action. Threat of geolocation is avoided if you choose a goal that doesn't ever lead to this type of investigation. The threat is simply gone.

Remediated — A threat is remediated if the probability of it having an impact is reduced. Its root causes are identified, and alterations to the strategy can prevent it from occurring in the first place. The threat of arrest via analysis of SMS messages collected via dragnet surveillance can be remediated by the use of end-to-end encrypted messengers (even if the content of those messages is equally incriminating). The threat of identification via fingerprints is remediated by wearing durable gloves and scrubbing down all items taken to an action (because there will be a forensics team, but they will find no fingerprints).

Mitigated — A threat is mitigated if its impact is reduced. It may not be possible to avoid the threat or reduce its probability in any way, but the severity of its impact may nonetheless be reduced. A strict code of silence mitigates the impact of one person in a crew being arrested because that person may still face prison, but the wider impact to the other 6 individuals in the crew is reduced. Having a legal advisor on retainer mitigates the effects of arrest because someone qualified is certain to assist you during the investigation.

There is not a clear line between remediations and mitigations, but it is helpful to think of the two ways risk is countered. We can make it less likely to happen (remediated) and/or less impactful when it does (mitigated).

There is a fifth non-resolution to threats which is that they are ignored. A threat is ignored when it is not known either through ignorance or when the subjects of a threat model choose not to analyze it. For example, a member of a crew that feeds the homeless might call attention to the fact that fascists have been attacking the homeless and their defenders. If the rest of the crew says “that’s not a problem” without discussing it, this is ignoring the threat, not choosing to accept it. Accepting a threat leads to informed consent about the action. Ignoring a threat does not.

In order to identify countermeasures, take your ranked list of threats and use a mind map to brainstorm possible ways to resolve them. Some countermeasures will apply to multiple threats (see the example in Figure 3). Some threats will need multiple countermeasures to be stacked in order to be resolved. Some countermeasures may be redundant, and this can be a good thing.¹⁵ If you have a detailed threat library, you may be able to pull countermeasures from it. If not, save the ideas you brainstorm during the session and update the library afterwards. You may need to do research to validate that these countermeasures are sufficiently effective.

Drawing many lines between nodes on a mind map can lead to clutter. Another method is to label all your threats with markers like T1, T2, etc. or short labels like T-PHONE-HACK, T-FACIAL-RECOGNITION, etc. Do the same with your countermeasures like C1, C2, etc. or C-NO-PHONE, C-MASKS, etc. Make a table of your threats and simply list off which countermeasures apply to which threats (Table 3).

¹⁵This is known as “defense in depth” or the “Swiss cheese model.”

Figure 3: Threats and Countermeasures: Mind Map

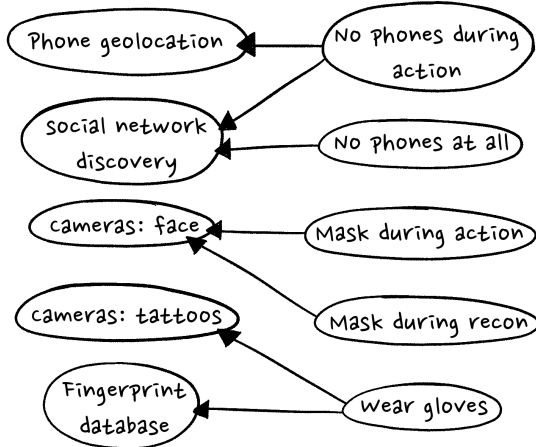


Table 3: Threats and Countermeasures: Table

Threat	Countermeasures
T1	C1, C3
T2	C2
T3	C4, C5, C6
T4	<i>none! needs work!</i>
T5	C2

When this part of the method is complete, you will have a complete threat model, and the next step is to use it to inform a security protocol. If you are unhappy with the model, you may need to iterate on it. While identifying countermeasures, you may have discovered that you have new goals to consider. While identifying threats, you may have found new adversaries to consider. Likewise, once you start devising your security protocol, you may find that the threat model has gaps that need to be filled in.

DEVISE A SECURITY PROTOCOL

By this stage, your threat model likely has many goals with many paths to them, and each goal and task may have many associated threats. You only need one strategy (set of paths) to reach your goal. Different strategies will lead to different security protocols, and each will have different associated risks. Some protocols may be too cumbersome to actually apply either by being too complex to execute or by

excessively impeding progress toward a goal. Some protocols might insufficiently address risk.

Pick a possible strategy that allows you to reach your goal, and copy the sub-goals and threats to a separate sheet of paper including recreating the connections between them. In the corner, write down two numbers that describe the risk of the strategy. The first is the highest risk number of all the threats (the maximum). The second is the sum of the risks of all the threats (the total).

$$\text{Risk}_{\max} = \max(T_1, T_2, \dots, T_n)$$

$$\text{Risk}_{\text{total}} = T_1 + T_2 + \dots + T_n$$

Next, discuss which countermeasures can be applied. Write down the countermeasures and link them to their respective threats. For each threat, recalculate its risk given the countermeasures applied to it. Under the two numbers in the corner, write down the new risk scores for this strategy. If you managed to reduce the total risk and maximum possible risk, the strategy may be acceptable.

Repeat this process with other possible strategies. If you feel that one strategy is less risky, but its risk number is higher than another, it doesn't mean your intuition is wrong. The numbers are in no way absolute. They are there to make you pause and think. If your intuition isn't in alignment with the numbers, this is something to investigate. Figure out *why* the numbers feel wrong. Maybe there's a huge threat whose countermeasures' effectivenesses were overestimated. Maybe it's the other way around and you assigned far too much impact to something. Adjust the scores on your strategies.

Once you have selected a strategy and the associated countermeasures, your security protocol is complete.

APPLICATION AND EVALUATION

Apply the protocol, and carry out your actions. Be mindful about whether everyone is following it. This requires discipline (for yourself) and trust (for others). If you don't have enough trust to know that everyone will speak up if they can't follow it, you may have insufficient trust to carry out the action. Or, perhaps you've planned this in to your threat model in which case minor deviations won't derail the project.

Over time, and especially after the action, evaluate whether the model seemed to match reality or if the protocol was effective. If you carry out many similar minor actions and keep getting disrupted while doing them, your threat model might not be addressing something. Reconvene to discuss the model when failures become apparent. Periodic discussions can be useful because it may be possible to slacken security, or it may be possible to add more countermeasures that were originally deemed to be too difficult to apply continuously. At a bare minimum, you alone or

your crew should redo your threat model once per year. New technologies, adversary tactics, or political landscapes will require reevaluation of the threat model.

It may be unsafe to share your exact threat model and security protocol with others, but you should try to engage in general discussions about security to gain insight into what other crews are doing. This can alert you to new trends in what your adversaries do or what your milieu does. You may find that crews you thought were safe to work with are not, or vice versa.

COMMENTS ON THE METHOD

After reading through the method, you might say to yourself “holy fuck mate, that is complex.” When it’s written down like this, it seems so, but in reality this is roughly the heuristic many of us use even if we can’t articulate it. Slowing down to name all the steps we might do in a handful of seconds makes it seem more complex than it really is, and once you’ve become familiar with common threats and countermeasures, threat modeling can be very quick and can require less use of numbers to assist with prioritization. I’ve seen crews made of veteran radicals devise new security protocols in as little 15 minutes if people are well-informed before arriving and tolerances for risk are roughly equal. Most of this time is just establishing whether everyone is on the same page, and once it is, quickly writing down the protocol is rather straight forward.

That said, threat modeling with this level of specificity and depth can be overkill for many scenarios. A typical process might be to only do cursory research into adversaries and then realize that other crews who have used similar strategies faced minimal repression. The crew might then decide to take standard¹⁶ security measures like not being overly talkative about it in public or on social media. This is completely fine.

Some actions need more focus, and this method is most applicable to those. However, this process requires practice. To completely nail it on the first try for a major action is unlikely. If you hope to threat model for something particularly interesting, you probably want to practice this method with your crew for smaller things to get habituated to using the method and to following an explicit security protocol.

EXAMPLES

To make the previously discussed process of threat modeling a bit more concrete, here are a few examples. One follows the method somewhat strictly, and this is not because everyone *must* follow it to the letter, but because it is an intentionally

¹⁶Standard relative to that context both spatially and temporally as well as relative to the “intensity” of the actions.

detailed example. The others do not for reasons that will be clear. As a reminder, while they are derived from actual cases, and while you may be able to apply elements to yourself, you should avoid slapping the resulting models on to your life without modification assuming they will pan out fortuitously.

As note on formatting, the people working out these scenarios would have lots of space and paper to use mind maps, but we're trying to cram this into an A5-sized zine, so all the iterations and notes aren't included. Abridged and somewhat final versions of their models are used to save space.

TYRE EXTINGUISHERS

SCENARIO

Some friends are part of a youth group for the local branch of their city's Green party. They are fed up with all the politicking and how little progress they have to show for all their time a meetings, conferences, and demonstrations. They've never taken direct action, but after seeing some social media posts about the Tyre Extinguishers,¹⁷ they decide to try that tactic. They sit down one evening to figure out how to do it.

THREAT MODELING AND SECURITY PROTOCOL

At first, the group has two goals: deflate SUV tyres, and definitely don't get caught by cops or the cars' owners (both are too scary). From the Tyre Extinguishers' website, the friends know they only need green lentils to deflate the tyres and flyers to get the message across. They brainstorm a flow chart of their goals and requirements (Figure 4). As they brainstorm, it becomes clear one of their goals isn't real. They don't care about deflating tyres. They care about discouraging people from driving large vehicles.

Once they have their goals and sub-goals, they think about who would stop them: the police and private individuals (either the owner of the vehicle or outraged neighbors). They start making a table of all capabilities these two adversaries might have (Table 4). They mostly use pop culture to understand how people might be caught, and this seems detailed enough for them. This seems fine until one friend points out that some printers leave barely visible yellow dots on pages to identify which device they were printed from.¹⁸ They do a little more research to see if there's anything else unknown like this.

Once they have their threats listed, they give each of them a risk score (Table 5). From this we can see that the biggest risks according to their model are:

¹⁷Call them libs if you want, but they've inspired people to take autonomous direct action instead of performatively getting arrested. Everyone has to start somewhere.

¹⁸These dots are called a Machine Identification Codes, and they've been around since the '80s.

Figure 4: Tyre Extinguishing Goals

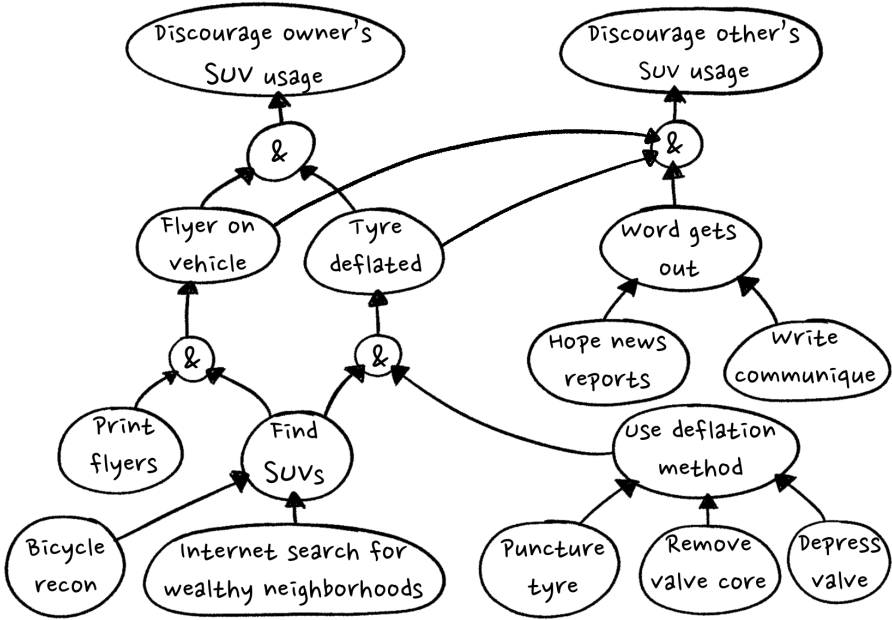


Table 4: Tyre Extinguisher Adversaries and Capabilities

Threat	Cops	Civilians
T-CCTV	Examine CCTV and car camera footage	Examine car camera footage
T-NOTICE	Notice people at night	Also notice people
T-SOCIAL	Look for relevant social media posts	Also look at social media
T-FINGER	Collect fingerprints	—
T-PHYS	—	Physically intervene
T-PHONE	Query mobile phone geo-database	—
T-PRINTER	Query printer dot database	—
T-EMAIL	Get email logs of who sent the commu- nique	—
T-WEB	Get internet traffic logs of who sent the communique	—

Table 5: Risk Scores

Threat	Cops			Civilians		
	Prob.	Imp.	Risk	Prob.	Imp.	Risk
T-CCTV	5	10	50	5	5	25
T-NOTICE	2	3	6	2	1	2
T-SOCIAL	2	3	6	2	3	6
T-FINGER	3	5	15	0	0	0
T-PHYS	1	10	10	3	10	30
T-PHONE	3	10	30	0	0	0
T-PRINTER	2	3	6	0	0	0
T-EMAIL	2	5	10	0	0	0
T-WEB	2	5	10	0	0	0

1. (50) T-CCTV/POLICE: cops catch them on CCTV
2. (30) T-PHYS/CIVILIAN: a civilian physically intervenes
3. (30) T-PHONE/POLICE: cops check phone records
4. (25) T-CCTV/CIVILIAN: a civilian checks their car's security camera and forwards it to the cops or publishes their face to social media.

Their maximum risk is 50 (the highest the scale goes), and their total risk after summing the cop/risk and civilian/risk columns is 206. They start looking at countermeasures for all of their threats. They begin organizing them into a table and matching them against the known threats (Table 6).

The groups proposes the following security protocol:

- Print flyers at their uni using their student association's printer that requires no credentials.
- Only quietly deflate tyres with lentils late at night when their owners and nosy neighbors won't be around.
- Wear masks to cover their faces when they're doing it.
- Get nondescript clothes from a secondhand shop that they only wear for these actions, and not dressing in all black as that draws more attention.
- Leave their mobiles at home.

Intentionally excluded from the protocol is a ban on talking about it with friends because while it's not an explicit goal, they want to inspire others to do the same, and they figure that hearsay is insufficient to get anyone to investigate them.

They then recreate the risk table to see if their countermeasures made a meaningful impact or not (Table 7).

Table 6: Tyre Extinguisher Countermeasures

Threat	Countermeasure	Rationale
T-CCTV	C-MASK	Masks make their faces unidentifiable
	C-HAT	Head coverings make other features unidentifiable
	C-BORING	Boring clothes won't uniquely identify them
T-NOTICE	C-BORING	Boring clothes are inconspicuous
	C-NIGHT	Night provides decent cover of anonymity
T-SOCIAL	C-NO-SOCIAL	Avoiding social media means no evidence
T-FINGER	C-GLOVES	Wearing gloves means no fingerprints on the vehicles
T-PHYS	C-SILENCE	Being quiet means no one will come out of their home to investigate
	C-BORING	Being inconspicuous means they will only get noticed during the few seconds they're deflating the tyres
	C-NIGHT	People are less likely to even be awake to confront them
T-PHONE	C-NO-PHONE	No phones means no location data left behind
T-PRINTER	C-PUBLIC-PRINTER	A public printer has too many people using it to be as easily traceable
T-EMAIL	C-NO-EMAIL	No emails sent means no way to trace it back
T-WEB	C-TOR	Tor Browser for submitting the communique is anonymous enough

Table 7: Modified Risk Scores

Threat	Cops			Civilians		
	Prob.	Imp.	Risk	Prob.	Imp.	Risk
T-CCTV	5	3	15	5	2	10
T-NOTICE	1	3	3	2	1	2
T-SOCIAL	0	3	0	0	3	0
T-FINGER	3	0	0	0	0	0
T-PHYS	1	10	10	1	10	10
T-PHONE	0	10	0	0	0	0
T-PRINTER	2	1	1	0	0	0
T-EMAIL	0	5	0	0	0	0
T-WEB	2	5	10	0	0	0

Their new maximum risk is 15, and their new total risk 61. That's a 70% decrease in their maximum risk, and (coincidentally) just a hair over 70% decrease in their total risk. The group thinks their estimates are sound enough that these results are reasonable, and they decide this security protocol is sufficiently effective.

ANALYSIS

All of the elements of their security protocol will help them avoid getting caught, and depending on how much effort police put into investigating their acts as vandalism, this may be enough. Their protocol has places it could be improved with minimal effort.

They haven't considered that during the research phase they they should be cautious. If they do bike recon, they might want to do it in a way that marginally hides their identity but doesn't make them conspicuous. If they search online and use maps or street view, they should use Tor Browser to prevent having their IP addresses or cookies tie them to the locations where they carry out their actions. But realistically, these countermeasures are probably unnecessary.

They also haven't considered that travel to and from their targets might identify them. They might want to consider wearing one outfit to a park, swapping to their nondescript outfit, stashing their bags, and the continuing to their target. They also haven't considered that they should be masking on approach to their target. This might make them more conspicuous, but it prevents a CCTV down the street from capturing their faces and being used to identify them. Hearsay might actually be enough to get police to come knocking, but if their other countermeasures were successfully applied and they shut the fuck up when the police ask questions, they likely will evade trouble. Like with the caution on research, these too are probably unnecessary countermeasures.

One thing they haven't considered is how to reduce the impact of a physical confrontation. Someone mad about their car being "damaged" might get violent. Has the group agreed to stay and fight if someone gets grabbed? Should they carry pepper spray to get their assailants to back off? A common error in threat modeling is to only consider the "happy path" and develop countermeasures against threats when things are going according to plan. They should consider the "sad path" and develop countermeasures for likely scenarios once things start to go off the rails.

Similarly, they haven't agreed to a code of silence. If one person gets busted, will they rat out the others? Maybe. Maybe not. It wasn't discussed, and this will lead to tension if someone doesn't behave as others assume they should.

Assuming this is taking place in the so-called West, there is minimal repression against these minor acts of sabotage,¹⁹ and their countermeasures are probably sufficient to prevent them from being identified.

¹⁹With the exception of the US where there is a legitimate chance they will be shot by someone defending their "property" from "terrorism."

AN INFOSHOP UNDER THREAT

SCENARIO

There is a small anarchist infoshop called The Black Flag (TBF) that's run by a collective of the same name. They sell some books, hand out zines and stickers, and let people use the space for meetings and informational events. TBF was raided by the police on the accusation of distributing seditious material, though the investigation was halted and no charges were brought against them. TBF plans a meeting to figure out if and how it should continue operating, and more specifically what the risks of any decision might be.

THREAT MODELING AND SECURITY PROTOCOL

Members of the TBF start their threat modeling process by considering what the police are doing, not what their goals are (i.e., they are using adversary-oriented modeling). What they know is:

- The police were “tipped off” (but may have fabricated the existence of informant or the tip entirely) about illegal material at TBF.
- The police carried out a raid and seized material that may be found to be seditious in a court which would implicate the TBF members in criminal activity.
- Such a raid might happen again.
- Raids have scared off people who visited the shop.
- The shop is likely under surveillance (if it wasn't already).

TBF has had the nominal goals since its inception of spreading anarchist material and providing a space where anarchist ideas could develop and spread. Because each member has different ideas about what constitutes “anarchism,” there have been many different kinds of ideas, some of which conflict. In pursuit of these goals, and each individual's liberty, they have the implicit goals of not being raided, arrested, or firebombed by fash.

After a few rounds of debate, TBF has a few proposals to consider for how to keep operating. The collective seems to be splitting in to three factions based on what they think the best move is.²⁰

1. **Continuation Faction:** Stay open exactly as before.

- **Pros:** No capitulation to the State, continuing availability of materials, continuing support of all local groups.

²⁰Let's not assume they chose those names for themselves. I had to pick something moderately descriptive to make the narrative easier to follow.

- **Cons:** Possible increased risk of future raids, charges, and surveillance of individuals who use the space.
2. **Pragmatism Faction:** Stay open, but change the content of the books, zines, and events to things less likely to be labeled seditious.
 - **Pros:** Possible reduced risk of raids and charges, some material still available (better than nothing), better long-term strategy than letting TBF burn to make a point.
 - **Cons:** Letting the State dictate what material is “acceptable” within the context of anarchism without fully getting rid of the risk of raids.
 3. **Shutdown Faction:** Close the infoshop entirely.
 - **Pros:** Much lower chances of being arrested for sedition charges related to the existence of TBF.
 - **Cons:** Loss of a radical-owned space, harder to spread texts, harder to spread ideas.

During the debate, it became clear that there are conflicting goals. The Continuation Faction thinks that the most seditious material is that which is most worth spreading (high risk, but high reward). The Pragmatism Faction thinks some material is not particularly helpful in creating an anarchist world and is willing to sacrifice its availability in order to continue organizing the space for some forms of anarchism. The Shutdown Faction has two sub-factions: those who want to avoid all risk and those who think that the collective could continue to be effective while operating underground.

The meeting goes on and on, and after several more meetings it becomes clear that TBF has a few factors that make it impossible to keep operating as before:

- The Continuation Faction is unwavering in the desire to continue hosting the most seditious material regardless of risk.
- The Pragmatism Faction is unwilling to take on the risk of the Continuation Faction’s material at the possible expense of their liberty.
- The Shutdown Faction has one half that is only “fair-weather anarchists” and another that only used TBF because it was conveniently aligned with their goals, but they have no strong ties to it.

It appears there is an impasse because the group cannot align on their goals, let alone what risks they are willing to tolerate.

ANALYSIS

Each group was right for different reasons. Refusing to appease the State has its merits, and plenty of anarchists have defiantly published texts that landed them in

prison. Playing within the State's laws is practical as there's a lot less one can do to affect change from within prison²¹ than with the limited "freedoms" or "rights" we're afforded under so-called democracy. Shutting down and creating decentralized networks helps mitigate the ills of inflexible and outmoded organizations (but shutting down to avoid all risk just plain sucks lmao).

From this example, it might seem like there was a failure from the threat modeling exercise since there was no consensus on the goals, threats, and risks, and that no security protocol was produced. They didn't even get as far as modeling out risk scores or selecting a strategy. However, the threat modeling they attempted worked exactly as intended: it revealed that there were incompatibilities between their goals, and in this case the goals derived from their ideological lenses. The group should probably split up, though one of the two factions that want to stay open likely will "win" by retaining control of the space. Threat modeling is useful not just for identifying and managing risk, but as a tool to facilitate conversations about what we actually believe and what we think matters.

SOME CHEEKY PRANKS

SCENARIO

A city is in the midst of a prank war with many factions including the biggest faction of pranksters, the self-named anti-pranksters. The pranks have been escalating, and at the same time pranksters are snatched and held by anti-pranksters to prevent them from carrying out more pranks. Total bummer, dude.

Quinn, a prankster himself, decides to rally a crew to pull off an epic prank on a yet-to-be-decided notorious prankster. Through cautious conversations with friends and some knowledge of their ideological alignments and tolerances for risk, Quinn lets a few people know about a secretive meeting in hopes of assembling a small crew.

Quinn meets the potential crew at a park on a Friday evening then walks with them for a few blocks before pulling everyone into a loud pub and taking a booth in the back. Quinn lays out their rough plan of doing an epic prank using what might be described as near maximum security measures to protect themselves during the whole prank process from inception to years past completion. Because Quinn has preemptively filtered and vetted potential members, everyone is unsurprisingly on-board, and they begin threat modeling.

THREAT MODELING AND SECURITY PROTOCOL

The crew starts by roughly laying out phases for what they will need to do to maintain a high level of security against anti-pranksters and other prankster factions

²¹This isn't to say prisoners don't organize in anarchistic ways, but just that it's much more constrained.

during the planned prank. They give the prank four somewhat overlapping phases:

1. Planning: security for when and how they meet to exchange information and plan next steps.
2. Preparation: security for the information and resources gathering before the prank.
3. Execution: security for the prank itself.
4. Dissolution: security for the time after the prank has been executed.

The crew starts writing down ideas for how they'd get busted during each of the phases, things that would be generally true regardless of what strategy they eventually settle on. Their general list of threats are overwhelmingly related to surveillance and forensics (and in particular digital surveillance). Some of the surveillance threats may already exist for all of them because they are known to be prankster-affiliated themselves.

Some of the threats they list are:

- Tracking individuals' locations via mobile phones.
- Reading the crew's messages via hacked phones or computers.
- Direct audio/visual surveillance of the individuals' residences.
- Personal surveillance via agents on foot or in vans (and also CCTVs).
- Snitches and infiltrators ratting them out.

After writing this down, Quinn's strategy for selecting meeting spaces becomes obvious. Random loud places provide reasonable cover from being tailed, and they can't be bugged in advance. Since there could be the possibility of them being followed, they agree to leave well before future meeting times to give them opportunities to engage in anti-surveillance drills²² during their journey.

A second thing they note is that while there are some countermeasure that could be applied to the various kinds of digital surveillance, they choose operate as if they are in a cyber-denied environment (i.e., the risk of being hacked or tracked is assumed to be so high that they allow zero use of electronics) because they very well may be. This avoids the risk of intercepted communications, and it is far easier than, for example, very carefully procuring burner phones and never making a mistake while using them. With these combined, they choose to select the next meeting time and location at each current meeting so that this information is never made digital. They devise a codephrase they can use to alert each other for the need for covert discussions ("Fancy a pint?") in the event someone misses a meeting and needs to

²²"Drilling" is a surveillance industry term for actions carried out to detect if one is being actively surveilled.

be given the time and location of the next one. They devise a second codephrase that will abort the prank and its planning if anyone feels they have become too directly surveilled or they have become otherwise compromised.

Lastly, they place a ban on discussing or even hinting at the existence of the crew. It is not to be discussed, not even indirectly by telling friends that they can't make some social appointment because of "something secret." The crew doesn't exist, and no one should even suspect it does much less that any of its members are in fact members of it at all.

The security protocol for the planning phase is:

- Meetings have a predetermined location.
- For meetings, members are to leave their electronics on and at home or work.
- They should engage in anti-surveillance drills en route to the meeting point and arrive punctually.
- They should dress nondescript to avoid drawing attention.
- The locations are never repeated.
- There is codephrase for requesting a means of receiving in person the next meeting's details.
- There is a codephrase for aborting the prank.
- No discussing the crew nor even indirectly hinting at it should be done under any circumstances.

The crew moves on to discussing security for reconnaissance. Because they already have gathered data on various bastards pranksters, some intel already exists on the individuals' personal laptops and phones, and therefore accessing this is acceptable. For new research such as further information gathering about a location or mapping routes, they agree to only use Tails²³ over random unfamiliar public WiFi networks while leaving their phones and other electronics at home. If they have to do physical reconnaissance or surveillance, they will likewise leave electronics at home, engage in anti-surveillance drills, and wear nondescript clothing that partially hides their identity to the largest extent possible without becoming suspicious. No electronics in this case also refers to personal automobiles as modern vehicles often have mobile network connectivity for receiving software updates or for mapping. While they don't know what their actual prank will be, they assume there will be some sort of investigation into and a possibly an attempted counter prank, so to avoid leaving evidence, they agree to use clean procedures for all the items that are acquired to reduce DNA or fingerprints from being left on them. The crew groans at this because they now what a pain in the arse this is, but they know it's the right thing to do.

The security protocol for the preparation phase is:

²³The Amnesiac Incognito Live System, a USB stick with a small operating system that only uses Tor for all internet traffic.

- No electronics are to be used (including cars) except Tails from random WiFi networks for research.
- Nondescript, identity masking clothing is to be used for surveillance.
- Engage in anti-surveillance drills en route to and from candidate targets.
- Clean procedures are used to reduce the amount of forensic evidence on items used in the execution phase.

The crew can generally understand what security they would need to pull off the prank, so they decide to break for the evening so they can research who they would like to prank and what options are available. They make plans for the next meeting, and then head separate ways.

Later, the crew reconvenes after having collected research and is ready to threat model their actual operation. They have a list of several pranksters, their home addresses, and their addresses of their businesses. The crew looks at what sorts of pranks have been done before, and while some are low risk, they also seem like they might not sufficiently tarnish reputations, and again, the members of the crew joined up to do an epic prank. After weighing a few options, and guessing at their feasibility and outcomes, they decide to sneak into one of the prankster's houses and quickly redecorate it.

They debate when this should be done. At night there's low visibility, but less traffic to hinder pursuers as they escape and fewer crowds to disappear into. During the day, they'd be easier to spot, but also it's more likely the prankster would be away at work instead of at home sleeping. Security cameras seem to be high enough resolution and well-functioning in low light that night's advantages might be slipping away. That said, most pranks still happen at night, and so few pranksters get caught. The crew reasons that night is still the best time, or rather, very early morning. The cover of night also provides the advantages of being able to swap clothes quickly in some dark corner without drawing attention. Finally, to avoid CCTVs (private and otherwise) from catching their movements so obviously, they agree to do have a rally point where they swap into clean single-use attire they've gotten from second-hand shops before making the final leg of the journey to the target. After, they will split again and dispose of prank materials and their change of clothes. Lastly, since there's some chance they will get caught during the execution of the prank, they agree to purge their living or work spaces of anything tied to the prank.

Their protocol for the execution phase is:

- Clean their spaces of prank-related items before the prank.
- Do the prank at night.
- Converge on a rally point, and swap to clean single-use clothing.
- After the prank, diverge and dispose of clothing and other items.

The crew finally plans the dissolution phase. Because the anti-pranksters might come after them or the targeted pranksters might go for retributive pranks, they agree to keep separated and a low profile. They set a check-in protocol to see if anyone seems to think they are being investigated. The schedule is for 3, 7, 14, and 30 days after the prank after which they imagine the heat will have dropped significantly. Finally, because of possibility that one of them might some time down the road be pressured to snitch, they agree that other than the single meeting on day 3 after the prank to discuss what went well or what went poorly, they will never again discuss the prank with each other. It will only live on as a memory to warm their hearts and soothe their consciences that they have not stood idly by as the prank war raged on.

Their protocol for the dissolution phase is:

- Have a single meeting to discuss how the prank went.
- Have a code of silence where they never discuss it again.
- Have a fixed schedule of minimal check-ins to ensure no one is under investigation.

ANALYSIS

Well, this is a fictional scenario, so it's a little hard to say if it really would work or not. It's also not even detailed enough to cover *everything*. For example, what exactly are clean procedures? Perhaps that's something for another zine. Maybe they could have used burner phones or encrypted radios for the execution phase or even the planning phase, but that would have introduced a different kind of complexity and additional cost of acquiring these items. Good models often follow the KISS principle: keep it super simple. This reduces the chances for human error.

As long term pranksters, there's a good chance they've kept up to date on modern anti-prankery and that they've done similar pranks before. They meticulously walked through everything they would need to do and considered how they might make mistakes at each step and how those mistakes might lead them to getting caught. They haven't modeled out risk scores because there's so much variance that it doesn't matter, and the consequences of making a mistake could be so dire. The crew has gone to the near maximum amount of effort to reducing risk that is possible, so comparing a before/after score is irrelevant because there's little more they could conceivably do. For *everything* they've identified as a threat, they've addressed as fully as they can. In fact, they may have over-addressed some things, and its possible that their forced use of Tails was more than strictly necessary and likewise with their use of random locations and drilling en route to the meetings. Sometimes minimally beneficial countermeasures in addition to *actually beneficial* countermeasures can provide psychological comfort as long as they aren't a burden and don't interfere with the necessary ones.

Will this stop them from getting busted or pranked back? As the saying goes, it's possible to commit no mistakes and still lose. Their threat model and security protocol help them do the best they can, but they know it's not risk free. That is simply how the game works.

WHERE THREAT MODELING GOES WRONG

When applying any method of threat models and creating security protocols, there are certain recurrent classes of failures. Some of these are discussed in the following sections.

OVER-MODELING

Once one starts, it can be tempting to make the most accurate threat model possible. One can get lost in the details and begin obsessing over every possible contingency as they strive to reduce the risk to zero. This will never be possible, and there will always be some risk. This paralysis can be overcome by aiming for goals with minimal repression and then working one's way toward bigger goals. Stickering and minor shoplifting can habituate one to acting outside the law. Tagging, banner drops, or deflating SUV tyres teaches one to take covert action in the night. If you find yourself terrified by the perceived might of your adversaries, consider first engaging in actions that are known to be minimally repressed.

UNFEASIBLE PROTOCOLS

One can make an actionable threat model and devise an effective security protocol that would allow actions to be taken out with minimal repression. If the new protocol requires far more effort than one's current protocol, it is unlikely that the full protocol will be able to be implemented. Forgetfulness and old habits are genuine concerns, and attempting to fully alter one's behaviors all at once is generally not possible. There will be gaps and errors. Protocols may need to be slowly implemented either one piece at a time or with increasing complexity, and this may mean picking goals or tactics that have fewer threats to begin with. Getting people to meet in the dead of night at a precise place and time without phones might turn out poorly. It might be easier to practice meeting up to hang out using these strategies to ensure people are punctual and can navigate without online maps.

LACK OF FUTURE THINKING

Another failure is closing off certain future strategies—or creating intolerable amounts of risk—by picking security protocols that leave one open to current surveillance

or repression. Civil disobedience that leads to arrest may make it harder to take future action due to harsher penalties or one's biometric data being stored in police databases. One might decide there's no risk associated with talking about potential crimes because they'll never actually do them... until one day they realize that doing them is the needed course of action. A question we should all ask ourselves is whether our current security protocol will harm our future self if we ever want to go beyond what we currently do.

Similarly, a strategy might call for a relatively lax security protocol, but choosing to use only the minimum security for the assumed threats can create the problem that action is constrained. Bringing a phone to a demo that is expected to be tame might preclude you from taking radical action if the police get aggressive. When looking at a protocol, one should consider if there is a reasonable chance that they would want to take further action if a situation changes, and if taking such action would create a great deal of risk because of the countermeasures that were omitted.

MODELING INSTEAD OF ACTION

Threat modeling is a slow exercise when it is done thoroughly, but sometimes we have to act quickly. There may be a coup attempt or fascists might riot through the streets, and we will only be able to rely on standard operating procedures. There may not be time to model the situation, and even a quick mental risk assessment might leave you unwilling to act if you aim to keep your total risk at a comfortable low. The tides of risk are rising, and trying to retain your current level of perceived safety will eventually lead you to complete inaction. Not everything can be threat modeled, and there are times when the only thing to do is to act decisively, boldly, and quickly. Having established security practices can give you something easy to turn to when you need a quick and probably correct solution for a new situation.

BESPOKE THREAT MODELS

Anarchist crews often find themselves reinventing the wheel when it comes to theory, tactics, and organizational strategies. This is generally true, not just in regards to security. Sometimes this comes from naïveté in as much as one hasn't been exposed to the relevant ideas or texts. In other cases it comes from the arrogance of one thinking that they or their crew is so special that the ideas espoused by others couldn't possibly apply to their unique situation. They will go off on their own and try to derive practices and ideas from scratch. On discursive topics like "what is the nature of anarchism itself?" this might quickly lead to convergence toward established ideas. In something that is more technical like security which depends on understanding underlying tech or observing the actions of law enforcement agencies, starting from scratch tends to have very slow convergence on established and verifiable practices, if it even converges at all.

To the extent that it's possible, relying on bespoke threat models and security protocols should be avoided. While it's true that the specific threats for a given class of person in a given place and time might be rather different from others, there are still significant overlaps between them. For information technologies such as computers, phones, and the internet itself, there is a general uniformity in the threats faced globally. Often what's needed is to see what threats are *actually present* in one's area and then mapping known countermeasures to them. The answers are often straightforward, and convoluted countermeasures and security might sound fuckin' dope and mega spy-like, but often they are based on poor understandings of how police or technologies operate, and their complexity can become a point of pride. There is some notoriety when one holds and demonstrates arcane knowledge, and crews can feel superior to others for developing customized protocols that one else has.

As much as you can, use the knowledge of others to inform your threat model. Determining which models or threat libraries have accurate understandings of the world can be difficult, so the task on this is often verification. Verifying and synthesizing existing knowledge is far easier and leads to far simpler models and protocols than trying incept everything from the æther.

IGNORING FUTURE RISK

The world is becoming more dangerous, especially to radicals. The most basic rights to protest that the State oh-so-generously affords us are being chipped away, and minimally disruptive civil disobedience is being increasingly criminalized. New threats to radical action and organizing are appearing, the frequency with which existing threats are deployed is increasing, and the impact of nearly every threat is becoming more significant. It cannot be overstated that ignoring future risk by maximizing safety in the present is a recipe for disaster. This is not to say one should quote/unquote “burn the whole fucker down” tomorrow by taking wildly risky action, but that we are right now—in all likelihood—facing far less repression than our future selves in 10 or even 5 years. Building a “perfect” security protocol that slowly and carefully navigates treacherous waters might still leave you shipwrecked in the oncoming storm. Avoiding risk now by taking less frequent or less intense action is just deferring that risk to the future.

This risk avoidance can happen in a number of ways. One might avoid “big” things because of how highly criminalized they are, but just as well, we might bend our existing tactics to avoid minor repression. We might be slightly less open to outsiders in hopes of deterring infiltrators, but this harms both current and future capacities. We might avoid shows of solidarity with criminalized groups such as not showing solidarity with Kurdish movements to avoid being prosecuted as “supporting terrorist organizations” as has recently become the case in so-called Sweden. Avoiding risk is a privileged position, and solidarity means taking on some of the risk that is directed at marginalized groups. Threat modeling informs us of what

risk exists, but if we aim to minimize risk for ourselves, we've lost a key element of what makes anarchism a worthwhile ideology: altruism and mutual aid.

UNDERESTIMATING REPEATED RISK

Humans are not very good at statistics. Consider the following scenario.

A crew with 6 members carries out a type of action repeatedly. Each time they do it, each member has a 0.5% chance of getting caught for it. How many times can they do it before there's a 50% chance of someone getting caught?

Twenty-three.²⁴

Something with a tiny, almost negligible probability of impact done every other week would lead to 50/50 odds of at least one of them being caught in about a year. This doesn't mean they're "safe" to do it on Night 1 and that they will only get caught on Night 23. They might get caught on Night 1 or 100, just the odds of it happening "on exactly Night 1" or "never until Night 100" are both very low.

This example is simplistic, and you could argue that the more times someone does something, the better they get at doing it, but just as well, they could get sloppy and complacent, or cops could have accumulated evidence with which to bust them. Even all this talk of models we're doing is fraught because *we don't actually know what's going on*. Much of this is guess work. We don't know what the probability of impact actually is. In the above example, if the probability isn't 0.5% but 1%, then the crew can only do the action 12 times before they have an over 50% chance of one or more of them getting caught. Can anyone really accurately estimate the difference between probabilities of 0.5% and 1%? Probably not. Maybe the probability is something totally different like 0.1% or 3% (115 or 4 actions before 50/50 odds, respectively).

Estimating probability is *very* hard, and getting a gut feeling for repeated risk is very non-intuitive. It's probably generally helpful to assume that something has a rare probability of impact that's done repeatedly becomes almost certain probability of impact if it's expected there will be 10 or more occurrences.²⁵

REPERCUSSIONS FROM PEERS AND THE STATE

When organizing, there are unspoken rules about what is considered "correct" by a given milieu. This can be on topics like how we organize actions, what methods collectives might use for consensus, and what sort of communiques are issued when some new hot topic arises. We are constantly considering what actions to take—at all levels—based on the judgement of our peers. When called out—that is when

²⁴This is a trivial case of the binomial distribution where we ask "what are the odds it never happens?" The formula that yields the answer is: $(1 - 0.005)^6 \cdot 23 \approx 0.5007$

²⁵Why 10? Honestly, we're making up numbers for so much, and we will never have enough data for a "scientifically" accurate model, but this seems reasonable. Ten is "some" but not "a ton."

we face repercussions—we can fall back on theory or say that we did everything according to protocol as if this exonerates us (and sometimes it does).

This is not the case with security. At the end of the day, what matters is whether our goals were reached and if we didn't get caught. If we follow a security protocol that is widely used within our milieu, if we get the nods of approval from our peers, and then we nevertheless end up in prison, the fact that our peers positively appraised our actions is meaningless.

A security protocol is not there to deflect criticism. It is not there to appease the wishes of your peers. It exists to keep you (relatively) safe while you carry out actions that lead you toward your goals. This sounds obvious, but it truly can be a paradigm shift for many. Someone might get busted and decry that “it's not fair” that they got caught because they “did everything right.” It doesn't matter if they used the protocol everyone else did. They got caught, which means it may have been insufficiently secure to actually protect them or that the action was inseparable from a large amount of risk. And again, what is considered “right” by your peers versus what is actually going to create security may not be as strongly correlated as you might think.

Your security protocol isn't about creating an image that others approve of. It's about genuine security. Do not forget this.

CLOSING REMARKS

When one learns to draw, they don't immediately produce life-like illustrations. What they make are misshapen heads or oddly proportioned critters over a messy background. With time and practice, they learn to more accurately represent their subjects. The shapes are more representative, and detail is added where needed and removed where unneeded to help shift focus.

When learning to threat model, your model might be clunky and simple, but as you practice, it will become more comprehensive and realistic. It's also okay to never become particularly elegant. As anyone who's played Pictionary knows, a quick drawing with only the most relevant details thrown in is enough to win a round. Sometimes people who can only make shitty sketches beat out artists who get hung up on complete representation. A simple threat model that identifies a few highly relevant things can be superior to a monstrously complex web of every possible contingency.

Threat modeling is an iterative process. Maybe there's only a single iteration and you call it good enough, or it's something you continuously come back to. Others' threat models can inform your own, especially through the intelligence they've gathered on shared adversaries. You might refine your threat model periodically or after any observed deficiencies. Your crew might need to revisit theirs if they have an escalation in activities or when a new member joins and brings their own insights or questions. A static threat model is itself a threat because as it becomes out

of date, it loses efficacy. Because the process is iterative, it's better to get a working model together and apply it in the real world and iterate than to spend months trying to create a "perfect" model.

There are a few extremes you should be wary of when discussing threat modeling and security in general. Security maximalists assume that without a robust threat model and airtight security protocol, any action will lead to your arrest or imprisonment. They tend to assume that all adversaries are or will be interested in you and will eventually act using their maximum capabilities. Being around maximalists can be anxiety inducing because no matter how much one does toward their security, it's never enough, and their paranoia tends to stifle action. Conversely, there are security minimalists who claim that everyone is too paranoid and that even going through threat modeling is a ridiculous exercise. They assume the most anyone will investigate them is only the local police who are too stupid to figure out what anyone is up to. Security nihilists believe—like the maximalists—that our adversaries are inhumanely powerful, but instead of trying to win the security arms race against them, the nihilists say that no amount of security can overcome the threat, so why bother?

I make a point of not knowing what activities anyone is actually doing, so I can't say for certain, but with some anecdotal data and intuition, it seems that security maximalists, minimalists, and nihilists aren't actually doing much of consequence. Those who do like to get spicy tend to have a much more nuanced take on security. Which is to say, regard extreme opinions with due caution.

Perhaps most importantly, threat modeling and the generated security protocols are only useful if they can actually be applied in real life and **if they help you achieve your goals**. A security protocol that is too cumbersome to actually be executed is useless no matter how well designed it is. Partial execution of the protocol in pursuit of the goals can create more risk than you intended. Similarly, a well-designed and easy-to-apply protocol that makes it near impossible to actually reach your goals indicates either a design or strategic flaw, or it indicates that your goals and your tolerance for risk are not compatible. This may mean restarting the threat modeling process, changing your goals, or finding ways to habituate yourself to risk and build up a tolerance.

Hopefully by reading this you can see why structured processes for evaluating risk exist. The world of secrecy and insurgency in which we operate is complex. We are constantly shrouded in the fog of war. Threat modeling can help part this shroud and give us a little more insight so that our actions can be done with confidence. We may still get caught, and there is no world in which all of us avoid repression. The stakes are high, but applying a bit of knowledge shifts the odds in our favor.

We ask our comrades "is it okay to use mobile phones?" and they respond "it depends on your threat model" before listing off scenarios that might endanger us. From these interactions, we have some idea of how a threat model is used, but how they're created is usually less clear. We shouldn't fully entrust our safety to others, and so we need to learn to manage the risks we ourselves face. This zine covers the methods you can use to threat model on your own and how these explicit steps can make you and those around you more safe. Our security is only as good as our models are accurate, so let's sit down and really think about how we can keep fascists and cops at bay as we move towards a liberated world.

