Handout 1: Accidents

Earlier, we discussed the question of whether technologies make our lives safer or more dangerous. We looked at this question in general and in relations of weapons.

We now turn to a more specific discussion of the **nature of accidents** in **high-risk technologies**. In doing so, we will consider how various *systems* of technologies expose us to catastrophic risk, specifically nuclear reactors (but the lesson here can be applied to any system of technology), and why preventing such catastrophes may be impossible. The goal here will be to understand a *certain type of risk* that technologies expose us to in the hopes that an increased understanding will allow us to reduce risk.

But, another goal is philosophical. If technologies expose us to unavoidable dangers, we are faced with the difficult question of whether the *benefits awarded by the technology are worth the risks*. How do we calculate the pleasures of cheap electricity or transportation against the horrors of radiation poisoning or death? In addition, we will consider what it means to say that a risk is *impossible to predict*, what this says about human beings as knowers and the relationship between technology and science.

1. Redundancies

Q1: How can we avoid error, decrease the number and seriousness of accidents?

It was suggested earlier that we might overcome certain risks through a variety of different means:

(i) rigorous testing
(ii) better training for personal (operator training)
(iii) developing safer designs
(iv) more quality control, routine safety inspections
(v) more regulation concerning safety.

Another way to avoid accidents is to build in redundancies. A **redundancy** is a safety device or process that aims to avoid some negative effect (accident) by (i) preventing negative impact if some process fails (passive redundancy) or (ii) minimizing or eliminating negative performance by allowing another "back-up" process to take over (active redundancy).[1]

Example 1: Giving a key to a friend (Perrow, p.6).
Goal: You want to get into your apartment. **Process #1:** You take your keys with you when you leave the apartment. **Result**: *You forget your keys!* **Process #2 (the Redundancy):** You go over to your friend's to get a spare key. **Result**: *He's there. Now you can get in your apartment.*

Process #2 is "redundant" because it performs the same result (achieves the same goal) as Process #1. But, notice that Process #2 isn't pointless: it allows you to achieve your goal if your primary process fails. This is important because the redundancy allows you to avoid the **bad consequences** if process #1 fails. In other words, it allows you to reduce your risk.

[1] (1) **Active redundancies** aim to minimize or avoid decreases in performance by substituting one process with another. These can be thought of as **back-up processes.** (2) **Passive redundancies** reduce *negative impact* of accidents or failure. These redundancies are commonly associated with **fail-safes (or fail-secures).**

Example 2: Building a Second Stairwell!
Goal: You want to get out of your apartment building. You live on the 5th floor. **Process #1:** You walk down the primary stairwell. **Result**: *There is a large fire in the 4th floor stairwell.* **Process #2 (the Redundancy):** You exit by the secondary stairwell located on the other side of the building. **Result**: *Now you can exit your apartment*.

Notice again that process #2 isn't necessary if process #1 works. No one may use the 2nd stairwell in normal circumstances. And, notice that the redundancy allows us to **avoid disaster** and it **reduces the risk of living on the 5th floor** of a building.

Other Examples

Example 1: Power distribution. There isn't a single wire connecting all of the houses together. If one wire goes down, current is redistributed through the other wires.

Example 2: Airplanes. Computers, hydraulic and propulsion systems are all duplicated (or triplicated) so that if the hydraulic system goes out, there is a back-up process that takes over.

Example 3: Handbrake/Parking Break. If your main brakes go out on your car, you can pull the parking brake.

Example 4: *Lawnmowers* have a lever that must be held closed by the operator's hand. If the lever is released, then the mower's blades stop. It is a mechanism built into the design to avoid danger, requires the mower to be held in a safe position at all times or shuts off.

Example 5: Certain trucks or trains have **air brakes**. These brakes make use of air pressure to keep the brakes in the off position (instead of using power to use the brakes). If there were damage to brake line, then the brakes are immediately applied.

Example 6: Bridges are built with extra cables and more cable strength than needed. If one cable brakes or degrades, the bridge does not collapse.

CDQ: Think of another redundancy. Articulate it by filling out the table below

Example 3:
Goal: **Process #1:** **Result**: **Process #2 (the Redundancy):** **Result**:

CDQ: Many people could reduce the risk in their lives by adding a few redundancies. Take a moment to create a redundancy for something in your life.

Example 4:
Goal: **Process #1:** **Result**: **Process #2 (the Redundancy):** **Result**:

Note that redundancies are built into our lives with the expectation that accidents happen! That is, we *anticipate* that we might forget our keys or that a fire might break out in a stairwell. A redundancy thus involves an ability to **anticipate the future events**. If a fire broke out in the 4th floor stairwell and no secondary stairwell existed, we could, of course, react in a variety of different ways, e.g. run through the flames, hide in the apartment as smoke fills the room and we pass out from smoke inhalation, leap from a 5th floor window to the ground. But, these actions are **not redundancies**. They are instead **reactions**. They

don't involve any form of **deliberation or anticipation** of future events. They don't involve putting processes in place to avoid an accident in another process fails. They are just responses to an accident occurring.

2. Two Questions

Given that building redundancies is key to avoiding accidents, two questions emerge:

> **Q2:** With a fairly high degree of forethought and inquiry, can we anticipate *all* of the *serious* accidents that might result from using technology?
> **Q3:** What is the *primary cause* of accidents?

Let's consider two answers to Q2.

> **A1 (Optimism about Risk)**: Yes. We **can** anticipate ***all*** of the *serious* accidents that might result from using technology. Human beings can work toward preventing dangers associated with accidents and ultimately we can prevent all serious accidents that result from our use of technology.
> **A2 (Pessimism about Risk)**: No. We **cannot** anticipate ***all*** of the *serious* accidents that might result from using technology. Human beings can work toward preventing dangers associated with accidents but these can never be fully prevented.

Notice that A1 and A2 seem to imply certain things about human beings as knowers, the complexity of the world, and the relationship between technology and science.

> **A1** implies that human beings are capable of knowing nearly everything, that the world isn't so complex that it is beyond our practical capacity to understand and anticipate future events, that reach of technology does not extend beyond the reach of science.
> **A2** implies that human beings have a limited ability to know/anticipate events in the world, that the world is so complex that it is beyond our practical capacity to anticipate *some important* future events, and that the reach of technology extends beyond the reach of science.

3. Perrow's Answer: Normal Accidents

One of the key claims of Perrow's *Normal Accidents* is this:

> There are certain high-risk technologies that "no matter how effective conventional safety devices are, there is a form of accident that is inevitable" (p.3) as these technologies expose us to catastrophic-like risk, we ought to consider either (i) abandoning these technologies or (ii) trying to understand them better to minimize risk.

The idea here is that a better understanding of risk will allow us to avoid simply putting the blame where it isn't, e.g. we might simply blame the operators.

PERROW'S ARGUMENT FROM NORMAL ACCIDENTS

P1	Most high-risk technologies have a special characteristic that make accidents in them inevitable or "normal."[2] This special characteristic concerns (i) the way that failures in them can interact and (ii) the way the system is connected.
P2	If we want to avoid catastrophe, our only choice is to (i) abandon some of these technologies or (ii) modify them in ways to diminish or reduce the harm it causes to human beings.

[2] These are normal *not in the sense that they happen frequently*; rather, they are "normal" in that they are inevitable or bound to occur.

IC Catastrophe is unavoidable.
C Therefore, we should abandon high-risk technologies or (if we can't) work toward modifying them or reducing their use.

But what is the argument for the unpredictability of accidents in high-risk technologies?

PERROW'S ARGUMENT FOR THE INEVITABILITY OF NORMAL ACCIDENTS

P1 Suppose there is a technology that (i) has a lot of components that interact in different ways and (ii) whose processes happen very fast.
P2 Now suppose that that system fails in a number *n* of different ways: Fy, Fx, …, F*n*.
P3 Designers can predict that Fy can happen & designers can predict that Fx can happen (we can build in redundancies), but it is almost impossible to predict what would happen if Fx&Fy, …, & F*n* happen (we can't build in redundancies for *all* of these accidents and the way they interact).
IC Therefore, the system is bound to produce some unpredicted accident.

O1: But we can see that accidents are occurring and make adjustment on the spot!
R1: Yes, but only if the technology has a lot of slack. That is, only if the processes of the technology can be turned off or don't happen so fast that we cannot react in time.

4. Interactive Complexity: An Everyday Example

In order to try and decide between **A1** and **A2** (and so answer **Q2**) but also *begin* to examine Perrow's argument, let's look at an extended example of an accident and consider **Q3**.

Example 5: Interactive Complexity

Goal: You want to get to class because you have an important final exam to take. You typically wake up 40 minutes before class, get ready, and then drive to work. It takes you 5 minutes to get ready, 5 minutes to drive to campus, and 5 minutes to walk to your class from the parking lot (leaving you 25 minutes to spare!)
(1) When you wake up, you realize you overslept 20 minutes, *but you can still make it to your class on time!*
There was a temporary power outage last night causing your alarm clock not to go off.
(2) You realize that you can make it to class and so you (in a rush) leave the apartment
***But* forget your keys and lock yourself out of your car & apartment**.
(3) You have a spare *car* key in your apartment just in case you lose your car key (a redundancy) but you've locked yourself out of your apartment! Luckily, you leave a spare *apartment* key at your friend's apartment (a redundancy) and so you decide to walk to your friend's apartment (she lives 5 minutes away)
***But* she's not there!**
(4) You walk over the bus station,
***But* the temporary power outage and the cold, snowy weather has caused a surge in the number of students needing to ride the bus and so the bus is full. In addition, there is a new university party holiday that begins in the morning so the bus is full of people trying to get to campus to party!**
(5) You try to call a friend who has a car to get a ride,
***But* she doesn't pick up (she might be in class or partying).**
(6) You try to call a cab
***But* no cabs are available because a number of students missed the bus had the same idea.**
(7) You knock on your neighbor's door and ask for a ride to the landlord. Your neighbor gives you a ride to his place, you pick up the spare key, return to your apartment, try to open the door,
***But* the key is not the right key (it is to another apartment)**
(8) You return to the landlord, he gives you a key. You return to your apartment (but your landlord cuts costs on the material of the spare keys) and
***But* the key breaks off in the door.**
(9) Finally, your neighbor says to you that she will take you to class (how nice!).

Result / Accident: You show up 5 minutes before your exam ends. Your teacher will not let you make it up, so you answer the questions quickly. And get an F for the exam and an F for the course.

One of the key ideas behind Perrow's argument is that some unpredictable accidents in technological systems often begins with failures that are

(i) *trivial* by themselves
(ii) to *some degree expected to occur in isolation* (and so we build redundancies),

But what we don't (and perhaps) can't anticipate and build in redundancies for are *all of the failures happening at once* and *interacting* in the way that they did. It is only when all of these failures happen and interact that the accident becomes **serious**.

First, let's look at some key events and think about how *unusual they are* and how *bad it is if it happens independent of the rest of the events*

	Event	**How Serious?**
1	Someone's alarm not going off	
2	a temporary power outage	
3	someone rushing to get to class	
4	someone locking themselves out of their car & apartment	
5	your friend not being at her apartment at the exact moment you need them to be there	
6	a bus being full	
7	no cabs being available	
8	your friend not picking up the phone when s/he is in class	
9	your landlord being busy giving you the wrong key	
10	your landlord trying to cut costs on supplies	
11	showing up very late to class and your instructor not letting you take an exam	

Next, let's consider whether certain events *depend* upon each other (and/or are frequently connected) or if they are *independent* (and/or if these events don't usually happen together) of each other

Event	**Linked?**	**Not Linked?**
power outage and alarm clocks not working		
many students *not* waking up on time during finals and a bus being full		
the buses being full and there being no cabs		
your friend being in class and not picking up the phone		
your landlord being busy and giving you the wrong key		
showing up very late to class and your instructor not letting you take an exam		

Finally, let's consider what the *primary cause* of your getting an F. Perrow's claim is that the *cause* of the accidents is multiple failures happening in a complex system and interacting in a way that we could not expect. That is, the claim is that the **cause of the error is the complexity and fast interaction of the technological system**.

Cause	**Example**	
Human Error?	Forgetting your keys *caused* you to fail the class.	*No! It might have been the first thing in the process but* ***normally, forgetting your keys***

		wouldn't be a problem; *you built in redundancies and explored other pathways, but these didn't work!*
Mechanical Failure?	The key breaking in the door *caused* you to fail the class.	*No! Normally, the key breaking wouldn't have been such a big deal. You would just notify the landlord to get it fixed and then take the bus to the exam.*
Environment?	The cause is the inconvenient bus schedule.	?
Design Flaw?	The cause is due to the stupid design of apartment door and the car. The fact that we live in a world where we need keys is the cause of the error.	?
Inappropriate Procedure	The cause is the *order* or *sequence* in which things are done.	?

We might say that **none** of the above can be said to be the **primary** cause of the accident. You prepared for many of the different possible events above, but what you did not predict was how all of the problems to occur and the interconnection between the different events. That is, you might be able to anticipate what would happen if you don't wake up on time by saying to yourself *oh! I will get my spare key from my friend, or oh! I will take the bus, or oh! I will take a cab*, but what you don't prepare for and didn't anticipate is

> that your friend's schedule might be different during exam week and that the power outage might cause a spike in the number of people riding the bus, which also leads to a spike in the number of people taking cabs, which affects your landlord who is dealing with things like snow removal and increased legislation from the borough and so has had to cut costs on materials and employees who have inadvertently mixed the spare keys up.

In other words, the accident is the result of the ***interaction*** of several different components and so what caused you to get an F was the ***complex interaction*** of the many events.

In thinking about an answer to **Q3**, let's consider several points.

Point #1. All of the above items are **trivial by themselves**. Independently, each of these can be expected. **However**, when **all of them of the above occur and interact together (coupled), and then** failure occurs.
Point #2. *We may want to blame people for mistakes, but how could the person have anticipated so many things would go wrong!* Especially when applied to highly complex systems, certain accidents seem unpredictable.
Point #3: Large accidents can have small beginnings. *Consider that failing a course (large accident) began with a temporary power outage*. We will look at how similar things happen with technology

So, let's reconsider **A1** and **A2**.

> **A1 (Optimism about Risk)**: Yes. We **can** anticipate ***all*** of the *serious* accidents that might result from using technology. Human beings can work toward preventing dangers associated with accidents and ultimately we can prevent all serious accidents that result from our use of technology.
> **A2 (Pessimism about Risk)**: No. We **cannot** anticipate ***all*** of the *serious* accidents that might result from using technology. Human beings can work toward preventing dangers associated with accidents but these can never be fully prevented.

CDQ: With a fairly high degree of forethought and inquiry, can we anticipate *all* of the *serious* accidents that might result from using technology? Give at least two reasons for your view *and* explain what your answer implies about the nature of human beings and their interaction with technologies.