Lessons from the Development of the Atomic Bomb

Toby Ord | November 2022

Centre for the Governance of AI

Lessons from the development of the atomic bomb

Toby Ord*

The creation of the atomic bomb is one of the most famous and well-studied examples of developing a transformative technology — one that changes the shape of human affairs. Teams of scientists and engineers in many different countries knew in advance that the technology may have tremendous implications for the world and strived to turn the idea into reality. The history of their projects provides many insights that may be useful as scientists and engineers today strive to develop new transformative technologies, such as artificial intelligence, synthetic biology, or nanotechnology.

There is much we don't know about the future development of these technologies. This makes it much more difficult to reason about the strategic landscape that surrounds them. Which, in turn, makes it more difficult to help make sure the development is safe and beneficial for humanity. It is thus very useful to have a case study of developing a transformative technology.

The making of the atomic bomb provides such a reference case. Its importance to world affairs was known from the start, and so it shows how the key players planned their routes through the challenges to reach this goal. It has been very well studied, giving us detailed information about which gambits succeeded or failed. And it was something of a natural experiment, being attempted by all five major powers in the Second World War. This allows us to see how things played out differently each time, getting us better information about which obstacles were really the hardest and which ingredients were required for success.

Of course there is no reason to think that things will play out quite the same way next time. For one thing, there are major dis-analogies. The world is not on the brink of total war and these new technologies are not weapons, so we hopefully will not have the same arms-race dynamic. Instead of competing national programmes, we may well have more companies involved, and even open academic efforts. And there are probably many more differences in both the challenge and the global milieu in which it is situated. Even if there were not, the process of developing a world-changing technology is complex and stochastic, with no guarantee it would turn out the same way twice.

So one should treat the development of the atomic bomb not as a map to one's destination, but as a detailed account of another traveller's journey in a nearby land. Something that provides valuable hints to important dangers or strategies we might not have considered, and which we neglect at our own peril.

---

* Future of Humanity Institute, University of Oxford.

I am very grateful for the help of Daniel Dewey, Katja Grace, Pablo Stafforini, Ian Goodfellow, and everyone at FHI, for their thoughtful comments and discussion.

In this report, I have attempted to summarise all the most important aspects of the development of atomic weapons and to draw out a number of important insights for the development of similarly important technologies.

- Section 1 — Lays out the key pre-requisites for making an atomic bomb, including scientific, engineering, and political challenges. It shows how these related to each other and gives numerous examples of how different nations failed to achieve them.
- Section 2 — Looks at the history of the Manhattan Project to show just how much labour and money were required to design and build the bomb, taking care to put this into context.
- Section 3 — Examines the role of secrecy in the different national programmes, showing when it worked, when it failed, and when it backfired.
- Section 4 — Examines the role of spying, focusing on the dramatic success of Soviet spying on the US and UK programmes.
- Section 5 — Explores the ability of scientists to provide decision makers with useful estimates of the cost and effects of an atomic bomb.
- Section 6 — Covers the decision-making in the US about whether and how to use the bomb on Japan.
- Section 7 — Describes the actual effects of the atomic bombs dropped on Japan.
- Section 8 — Details the efforts of the atomic scientists to control the development and use of the technology they were creating.
- Section 9 — Explores the profound impacts that individual actors had on the development of atomic weapons.
- Section 10 — Explores how scientists managed the potential existential risk of nuclear weapons igniting the atmosphere.
- Section 11 — A collection of additional brief insights on the process.

Prerequisites for the Atomic Bomb

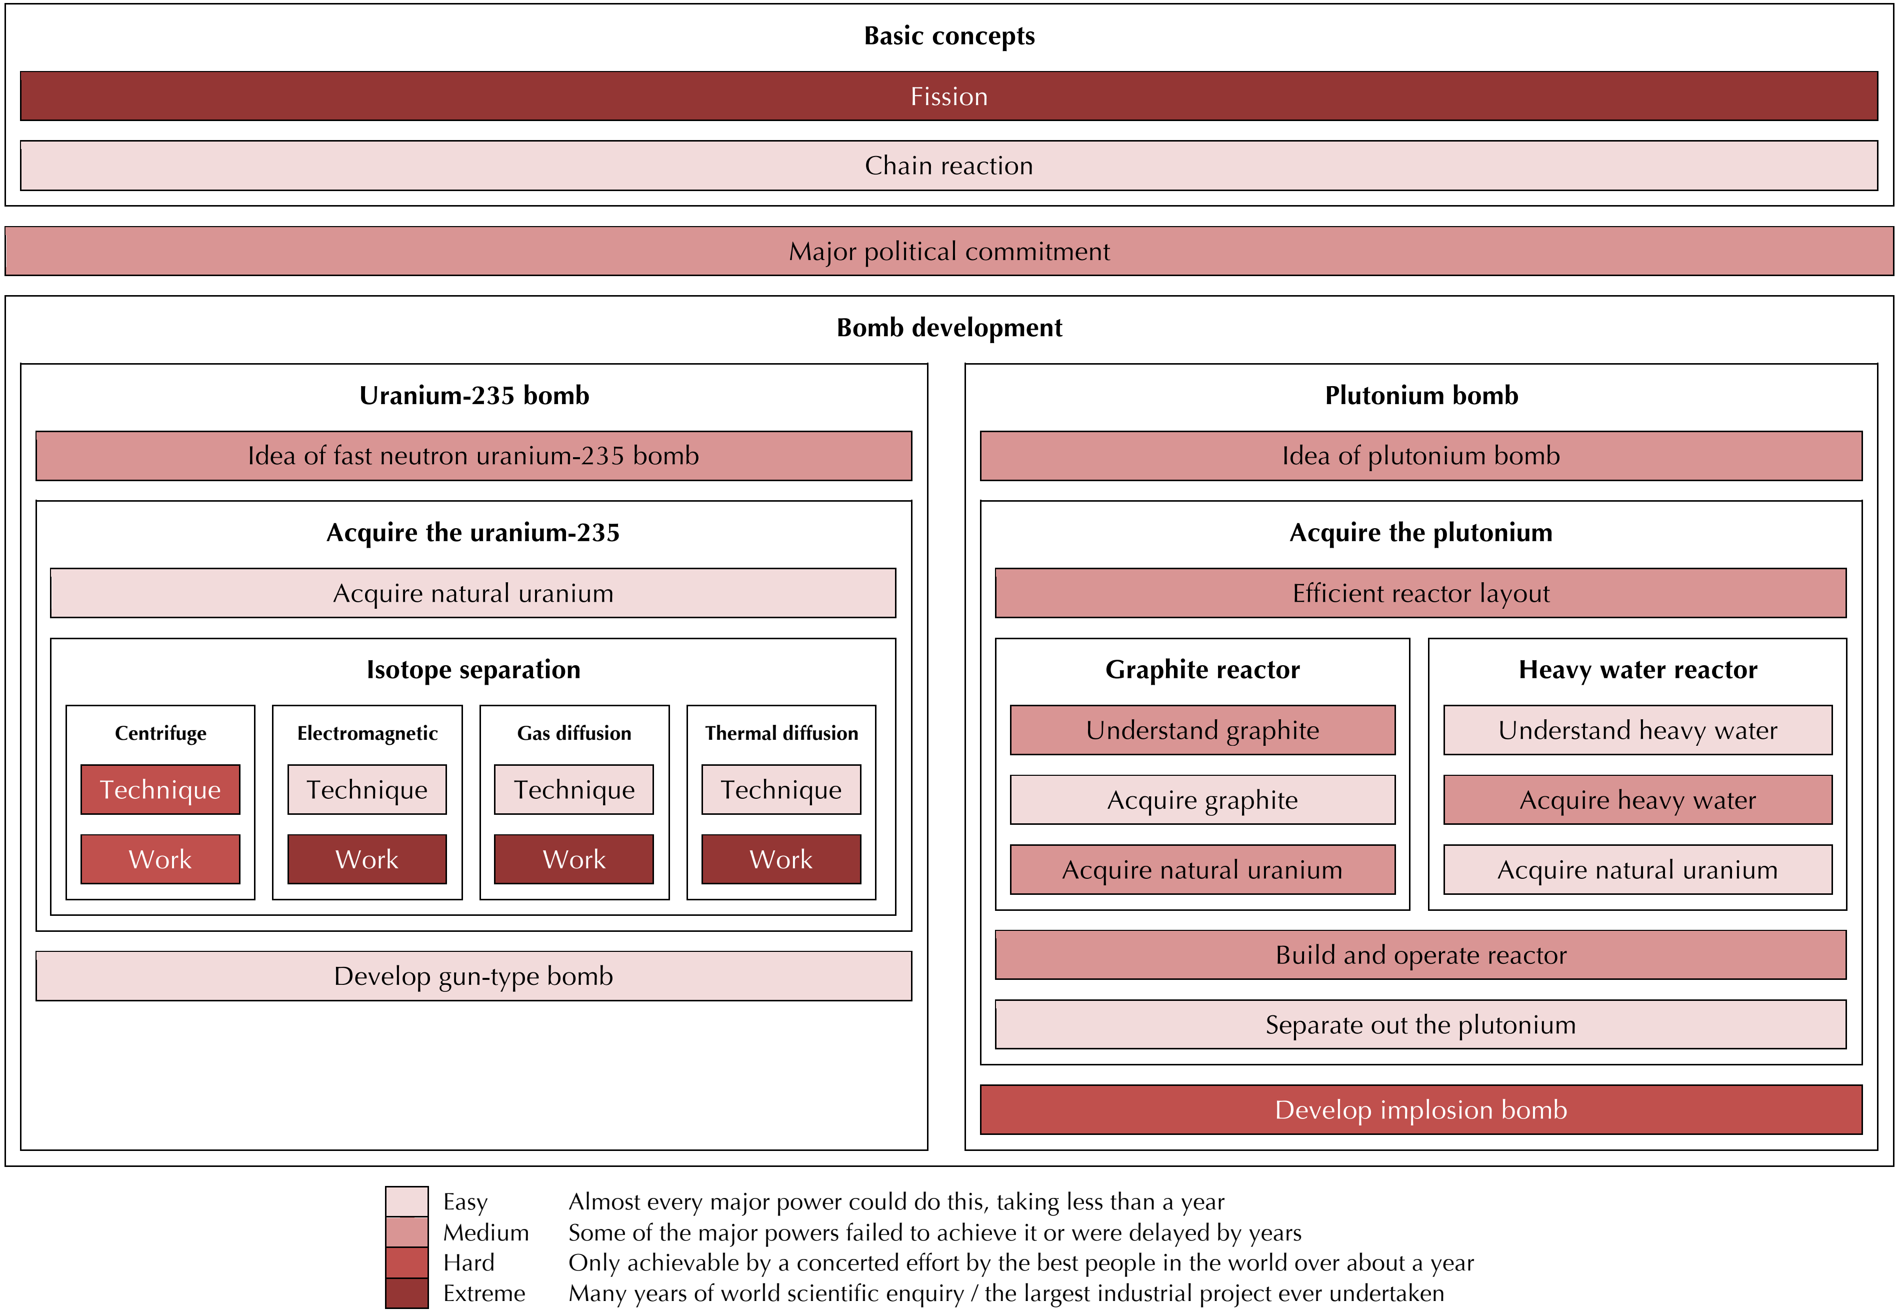

1. Pre-requisites for the atomic bomb

Developing the atomic bomb involved many important steps. These included not only scientific and engineering breakthroughs, but establishing political support, acquiring strategic resources, and a truly vast amount of more mundane work. The interconnections between these pre-requisites are also complex. In some areas, every step was required, while in others there were multiple routes to success.

Figure 1 shows the main pre-requisites for the building of an atomic bomb, their structure, and their difficulty. It provides a framework for understanding the roles of each of the key components, whose histories are related below. Similar diagrams could be useful for understanding the development of other historical or future transformative technologies.

Each box represents an important part of the process, the completion of which would help produce an atomic bomb. Some of these parts are further broken down into important sub-parts. This breakdown could have been continued indefinitely, but eventually the value of a more fine grained breakdown is outweighed by the increased complexity and one has to make a judgment call about where to stop. In this case, the coloured boxes represent parts that are not further broken down.

Parts can be broken down into a vertical or a horizontal set of sub-parts.[1] A vertical set means that to complete the part, all the sub-parts must be completed. A horizontal set means that to complete the part, at least one of the sub-parts must be completed. Vertical sets thus correspond to 'and', while horizontal sets correspond to 'or'. The nesting of these parts allows for complex logical relationships (indeed for all Boolean combinations). We can thus see from *Figure 1* that, at its most fundamental level, the bomb required the basic concepts, major political commitment, and bomb development. Moreover, the basic concepts required both fission and the chain reaction, while the bomb development required at least one of a uranium-235 bomb or a plutonium bomb. Each of which required many other things…

One important feature that is not captured by this diagram is the sequencing of the pre-requisites. While I have arranged the diagram to reflect a certain logical ordering (moving from the top down), this isn't formally implied by the diagram. A vertical set of sub-parts just means that all of these are required, but does not require that they are developed in order. As an example, all countries started designing nuclear reactors before having the idea for the plutonium bomb.

Another feature that is not captured is any more subtle relationships between different pre-requisites. For example, reactors were useful for understanding uranium well enough to make a uranium-235 bomb, but as there is no precise logical relationship between these, there is no connection on the diagram.

[1] The idea of structuring a diagram like this comes from Mark Miller.

A useful feature of this diagram is that all the minimal sets of pre-requisites for completing an atomic bomb correspond to vertical sequences of the coloured boxes reaching from the top of the diagram right to the bottom.[2]

As explained by the key, the colour of each box represents its difficulty, with more difficult pre-requisites being darker. I assigned these levels based on a combination of different sources of evidence: the time it took to successfully achieve them, the number of competing programs that failed to achieve them, and the specifics of the case. More detail can be found in the descriptions of each pre-requisite below.

I will now describe each pre-requisite in more detail. This will justify the structure of *Figure 1* and the difficulty levels assigned. It will also provide many more important details of the process that may be informative for the development of new transformative technologies.

Fission

Fission is when a neutron splits an atomic nucleus into smaller parts, releasing a portion of the energy within. It is the essential process that gives all nuclear weapons their power. In December 1938, a German chemist, Otto Hahn, discovered fission in uranium and published this breakthrough almost immediately (despite being concerned about the Nazi regime). Lise Meitner and Otto Frisch were the first physicists to be told and they finished the basic theory behind it within weeks. While physicists had caused fission in the lab before this (starting with Enrico Fermi's team in 1934), they had thought they were making new, larger elements and hadn't understood that they were splitting a nucleus into smaller pieces. The basic idea that uranium atoms were breaking apart into two or more large pieces under neutron bombardment was very simple and obvious in retrospect, but had defied discovery by the world physics and chemistry community for four years[3] and earned Hahn a Nobel prize. Its discovery precipitated the race to atomic weapons.

Chain reaction

While fission liberates nuclear energy from an individual atomic nucleus and creates a lot of energy per unit mass, one needs to scale this up by more than 20 orders of magnitude to make a useful proportion of the nuclei split. This can be done with an exponential chain reaction where the products of each atomic interaction create more than one additional reaction. The Hungarian physicist Leo Szilard came up with the idea of a neutron-based chain reaction in 1933 while based in the UK. He realised its implications, but was careful to keep it secret due to its potential military applications. The neutron chain reaction was independently discovered by the Soviet physicist Lev Landau in 1934. While Landau did not publish nor make much of the discovery prior to the news of fission, he did mention it to Peierls (who would later

[2] Other diagrams of the same type may not have this feature.

[3] The idea of fission had been proposed before this by Ida Noddack in 1934, but she had not provided any proof of the hypothesis and the paper was not influential.

go on to bring atomic weapons to national attention in the UK) and presumably Landau's familiarity with it helped the Soviets realise the importance of fission.[4]

When news of fission became public in January 1939, Szilard realised that it might produce enough neutrons to cause a chain reaction. He thought that Enrico Fermi or Frédéric Joliot-Curie may well realise the possibility for a chain reaction and decided to contact them to keep it secret, but was delayed by a bad cold.[5] While only two people had realised the possibility of a chain reaction in the prior six years, it was much more obvious after the discovery of fission. The theoretical possibility for a chain reaction was discovered independently by Fermi, by Joliot-Curie, and by Robert Oppenheimer all within about a week.[6] When Szilard recovered from his cold, he sought out Fermi and started an on-going struggle to get him to be quiet about it and subsequent discoveries.[7] He also contacted Joliot-Curie (and his team) to try to get them to keep quiet, but they decided to publish, with influential articles in *Nature* on 18 March and 22 April (the latter of which measured the number of neutrons released, showing that there were enough to enable a chain reaction). At this point, there was no way to contain the news.

A week later, on 29 April there was an article in the *New York Times* on the possibility of large fission explosions. On the same day there was a secret German meeting on nuclear fission followed immediately by a secret nuclear project in the war office.[8] Overall, it seems that given the other essential part of the puzzle (fission), the chain reaction was very easy to discover.

Get major political commitment

Building an atomic bomb required a vast scientific programme and an even vaster industrial programme, both of which would be impossible within a reasonable time frame without major political support. At its peak the US programme was employing a thousandth of the entire population and using about a two-thousandth of GDP.[9]

Germany and the UK very quickly achieved moderate levels of political commitment (in 1939 and 1940 respectively). In 1941, Japan also managed to achieve moderate political commitment, quickly scaling to about 100 researchers. The Soviet Union had an official nuclear programme from 1939, but was moved to much lower priority in 1941 during the German invasion and only reached a moderate level of political commitment four years later, in 1943, with about 20 scientists and 30 support personnel.

---

[4] (Rhodes 1995, p 32).

[5] (Rhodes 1986, pp 266–7).

[6] (Rhodes 1986, pp 271, 274–5).

[7] (Rhodes 1986, pp 280–1).

[8] (Rhodes 1986, p 296).

[9] See Section 2 for much more detail.

Despite getting support from President Franklin Roosevelt through the famous Einstein-Szilard letter of August 1939, the US was stuck with its nuclear programme in a badly run, resource-poor committee for three years (half the entire war). The UK got ahead of them on theory and sent a copy of their cutting edge MAUD report to the US. Eventually Mark Oliphant (an Australian working in the UK programme) flew to the US to find out why they hadn't responded to the MAUD report, and found that the head of the US programme (Lyman Briggs) had just filed it away and not shown it to the scientists. Oliphant directly told Ernest Lawrence about the recent British results and they both championed the case for pushing forward in Washington.[10] This led to major commitment (much bigger than any other nation), a large budget, and the appointment of extremely competent managers: Leslie Groves and Oppenheimer. By way of comparison, the US spent more than a hundred times as much in the last three years (once major commitment had been achieved) as it did in the three years before that.[11]

Bomb development

There were two main types of atomic bomb that could be made and they required very different techniques. The uranium-235 bomb was conceptually quite simple, but required an extraordinary amount of money and low-skilled labour — larger than any previous industrial project ever undertaken. The biggest challenges for the plutonium bomb were conceptual and technical challenges around implosion — getting the plutonium for the bombs was comparatively simple and cheap. The US did both in parallel and despite them being completely different processes with different difficulty, finished the first bomb of each within 3 days of each other. This was probably a testament to their strict timetabling (the bombs needed to be produced before the war was over) and their willingness to spend money to meet it. The US's uranium-235 bomb process ended up costing about three times as much as the plutonium bomb process and produced fewer bombs.

Acquire natural uranium

Naturally occurring uranium is a mixture of different isotopes, with the vast majority being uranium-238 and a mere 0.7% of uranium-235. Natural uranium was the basis for both types of bombs, with the uranium-235 being separated out in one case and uranium-238 being converted to plutonium in a reactor in the other case. The US uranium-235 bomb required about 10 tonnes of natural uranium per bomb. The US plutonium bomb programme used about 500 tonnes of natural uranium (producing several bombs).

Acquiring the uranium turned out to be relatively easy for most powers, but was a significant constraint for the USSR and Japan. When the US and USSR raced to control as much of fallen Germany as possible, taking German uranium was a top priority. The uranium the Soviets found hastened the start-up of their first

[10] (Rhodes 1986, p 372).

[11] $1.8 billion vs. $15 million, from (Hewlett & Anderson 1962, p 724). See *Figure 2* for more detail.

plutonium reactor by about a year.[12] The US and UK went to some effort to try to control world uranium supplies, controlling about 90% of the world's high grade ore by the end of the war.[13] But while extracted uranium was somewhat scarce during this time, it is sufficiently abundant in low-grade ore to make such control impossible in the long run.

Idea of fast neutron uranium-235 bomb

From the knowledge that fission chain reactions were possible, it was a very small step to consider a bomb using uranium-235. When designing a uranium-235 bomb, it was most natural to consider using slow neutrons as they are more readily absorbed. However, a slow neutron bomb would not work, as the heat of the reaction would separate the uranium atoms before the reaction could be completed, resulting in a fizzle rather than an explosion. The US programme did not realise this problem and did not consider the fast neutron uranium-235 approach for two years, at which point the UK programme told them about the problem and its solution.[14] The evidence suggests that the German programme also did not consider a fast neutron bomb.[15] The only other team that definitely worked out fast neutrons were needed was the Soviet programme.[16]

Isotope separation

Because uranium-235 is chemically identical to uranium-238, it is extremely costly and labour intensive to separate them, accounting for about 60% of the entire Manhattan Project budget. It requires so much work that in the early days many scientists were not sure whether to pursue the idea of the bomb due to the unprecedented industrial scale seeming impossible in practice. The major advantage of plutonium is that it can be chemically separated which is much easier. There were many proposed methods for isotope separation, and as with other choices, the US tried all of them in parallel to increase the odds that they would find a way that worked. Very late in the war, they realised that the processes had peak efficiencies for different levels of the enrichment process and chained them together in series, improving efficiency greatly.

Centrifuges

These were the most obvious possibility, but ran into practical troubles during development in the US, UK, and USSR, leading to them being abandoned for other approaches. Interestingly they were made to work after the war by small teams in the USSR and USA, becoming the most efficient method of isotope separation: about 50 times more efficient than other approaches. The innovations that enabled practical

---

[12] According to the head of the Soviet programme, Kurchatov (Rhodes 1995, pp 162–2).

[13] (Rhodes 1995, p 130).

[14] (Rhodes 1986, pp 373–4).

[15] (Kant 2002, p 15).

[16] (Holloway 1994, p 77).

centrifuges for isotope separation were relatively minor, all of them using pre-existing engineering techniques.[17] The main barrier to discovery during the Second World War seems to have been that the US and Soviet teams committed too quickly to a basic design and only considered minor tweaks rather than rethinking the main elements. It therefore seems that despite not being developed in this period, centrifuges could have been made to work.[18]

Electronic

This used a pre-existing technology (the mass spectrometer), so was the easiest to get working. It also worked well even at high purity levels. However it was the most expensive to use per unit of purification achieved (25% of the entire Manhattan Project budget).[19]

Gaseous diffusion

This was the workhorse of the US programme, but was still extremely expensive (27% of the entire Manhattan Project budget). An astounding 89% of this cost was in building the plant, rather than in running it.[20]

Thermal diffusion

This was not tried until late in the US project because it could not enrich uranium all the way to weapons grade (it thus shouldn't technically appear as an alternative approach in this diagram and is only here for completeness). The US started using it once they realised they could chain different enrichment approaches together, as it was efficient for the first stage of enrichment.

Develop gun-type bomb

The uranium-235 reaction could be harnessed with a relatively simple bomb. The main idea was to bring two subcritical masses of highly enriched uranium together sufficiently quickly. Effectively the bomb was designed like the barrel of a gun with an explosion at one end propelling a mass of uranium towards another mass waiting at the other end inside a tamper. Compared to the other parts of the bomb programme, this gun-type bomb was easy to design.

Idea of plutonium bomb

In 1939, the periodic table ended with element 92 — uranium. Despite the eagerness of scientists to create new man-made elements, they did not yet know how to make

---

[17] Being held up and spun electromagnetically, while pivoting on a needle, with loose bearing and dampening (Kemp 2012, p 282–3).

[18] There is an excellent discussion in (Kemp 2012).

[19] (Hewlett & Anderson 1962, p 723).

[20] (Hewlett & Anderson 1962, p 723).

them. Fermi thought he'd isolated them in 1934 and was even been awarded the Nobel Prize for this achievement in 1938. But when fission was discovered, everyone realised he had just been looking at a mix of smaller elements created by the fission. In 1940 in the US, a team led by Edwin McMillan discovered element 93 — neptunium — by bombarding uranium with neutrons and carefully isolating the neptunium from the fission fragments. In early 1941, the team (then led by Glenn Seaborg) managed to isolate element 94 — plutonium — which neptunium decayed into.

Plutonium-239 was suspected to have similar fission properties to uranium-235 and experiments by Seaborg showed this to be the case. However, it had the major advantage of being possible to purify it through chemical means instead of the much more laborious isotope separation. It was also simple enough to create that it should be produced in reasonable quantity in a nuclear reactor. It was thus quickly seen by the US as a candidate for a second type of bomb.

Every nation that considered plutonium bombs developed the idea before physically isolating any plutonium. The US worked out on theoretical grounds that isotope 239 of element 94 should work for a bomb and be produced in a reactor in May 1940.[21] On Szilard's advice, this was kept secret. Germany worked it out in July 1940 and then managed to isolate neptunium sometime before May 1941.[22] They searched for plutonium but couldn't find any due to lack of a strong enough neutron source.[23] The UK worked out the theory in early 1941.[24] The Soviets only found out about plutonium and its value for a bomb via espionage in the US in 1943, then found some more details by reading the article McMillan's group published in *Physical Review* 1940 on the creation of neptunium.[25]

Efficient reactor layout

Nuclear reactors involve a controlled fission chain reaction in natural uranium. This requires slowing neutrons with a moderator such as heavy water or graphite. All five great powers began theoretical and practical work on nuclear reactors before they knew they would also create the raw material for a plutonium bomb. Initially, they were pursued as a scientific test bed for understanding fission and as a potential source of power (for military or civilian purposes).

The moderator slows down the neutrons produced by fission and is needed because slower neutrons are more likely to trigger additional fission. The most obvious

[21] (Rhodes 1986, pp 346–7, 350).

[22] (Rhodes 1986, p 350), (Kant 2002, pp 7–8).

[23] (Kant 2002, p 7).

[24] (Rhodes 1986, p 350).

[25] The publication of the paper on neptunium attracted significant protest by scientists engaged in atomic research who saw its importance for this secret project. By the time of the discovery of plutonium, the Americans had realised the importance of this research and did not publish again. Interestingly, the Germans also published their discovery of neptunium (in *Die Naturwissenschaften* in 1942) (Kant 2002, p 8).

layout for a reactor is a uniform mixture of the uranium with the moderator. However this is inefficient: a rector structured like this cannot reach criticality without enriching the uranium beforehand. In July 1939 Fermi and Szilard quickly and independently realised that the uniform mixture was suboptimal and both considered a layered arrangement. A few days later, Szilard hit upon the optimal arrangement with spheres of uranium arranged in a lattice within the moderator.[26] Though they weren't yet sure of it, this arrangement is efficient enough to allow criticality with natural uranium.

In contrast, the Soviets only considered uniform mixtures, and in 1940 they determined that these required enriched uranium to work, seemingly putting reactors out of reach without isotope separation. They only corrected this after espionage in 1943 showed the US and UK were working on reactors with natural uranium, which spurred them to consider non-uniform mixtures and they almost immediately discovered the lattice arrangement.[27]

The German programme did most of its reactor experiments with layered arrangements — something that undoubtedly slowed their progress and may have contributed to the programme being deprioritised.[28] It was only in February 1945 that they first tried a lattice arrangement. This brought them very close to criticality, but came only a few weeks before they were overrun by Allied troops.[29]

Understand graphite reactor

Both the UK and Germany made incorrect measurements of the quality of graphite as a moderator (perhaps due to impurities in their samples) and they discarded it as a possibility. Its quality as a moderator shouldn't have been very hard to determine, but once an incorrect measurement was made, it was difficult to realise the need to try it again. It therefore ended up being a surprisingly important stumbling block on the path to an atomic bomb. The US later told the UK about their error, but the Germans never worked it out and had to rely on the other alternative of heavy water.

The Soviets correctly worked out how good a neutron moderator graphite would be (and even published this in *Physical Review* in 1940). But because they assumed a uniform mixture of uranium, they incorrectly thought a graphite reactor would require enriched uranium to work and so did not initially pursue it.

Acquire graphite

This was relatively easy compared to the other challenges, but required much purer graphite than usual, which required sophisticated industrial processes. The successful US programme (which mainly relied on graphite) used about 2,000 tonnes.

---

[26] (Rhodes 1986, pp 301–3).

[27] (Rhodes 1995, pp 39, 72–3).

[28] (Kant 2002, p 8).

[29] (Kant 2002, p 11).

Understand heavy water reactor

Everyone quickly worked out that it was a good neutron moderator. It was also easier to handle than graphite, and a heavy water reactor had the advantage of only requiring 1–2 tonnes of uranium instead of about 50 tonnes for a graphite reactor. This was very important for the Soviets who only had a couple of tonnes on hand in 1943. As with graphite (above), the Soviets did not consider non-uniform arrangements of the uranium, so they thought a reactor would require enriched uranium until espionage in 1943 put them on the right track.

Acquire heavy water

Production plants were very rare at the time, so it was hard to acquire. The Germans needed heavy water given their lack of understanding of graphite and noticed the world's main heavy water plant in Norway. When they invaded, they captured it and started using to help their nuclear programme. The Allies knew this was valuable so repeatedly sabotaged the plant and destroyed its stocks of heavy water. It is not clear whether enough heavy water made it to Germany for a viable plutonium bomb programme. The Soviets had very little heavy water and no production, but began building plants in 1943 once they realised its importance.

Separate out the plutonium

The primary reason for developing the plutonium bomb was that plutonium can be chemically separated from the raw uranium, which is much easier than the isotope separation required for uranium-235. However, it still involved substantial challenges. New chemical processes had to be developed and implemented on a massive industrial scale. The processing involved extreme radiation danger, which the US dealt with by using early closed-circuit television to create the world's first remote operated plant.

Develop implosion bomb

If it absorbs a second neutron, the desired plutonium-239 turns into plutonium-240. Impurities of this plutonium-240 (which are practically impossible to remove) make plutonium prone to premature chain reaction (and thus fizzling) unless the sub-critical masses are brought together faster than a gun-type bomb could achieve. Luckily, before this problem was discovered, the US had already considered the idea of a new type of bomb, which could bring the critical mass together faster by imploding it. They then realised this extremely challenging approach would be necessary.[30]

[30] It is interesting to consider what would have happened if this problem had been accurately foreseen earlier. It was such a challenge that it may well have made them abandon the plutonium approach (even though it would work out very well in the end). So counter-intuitively it might have been useful not to have known earlier.

The technique for implosion was exceptionally difficult, requiring many technical challenges and detailed understanding of the then-new idea of shaped charges. These charges acted as lenses for the explosion, focusing the explosive force inwards to the centre of the bomb, where the plutonium was located. Another very difficult aspect was the design of the 'initiator' at the very centre of the ball of plutonium.[31] The implosion bomb probably wouldn't have been possible in time without the involvement of three key experts: John von Neumann, Edward Teller, and George Kistiakowsky.

2. Required Resources

We are lucky to have good data on the resources used in the successful US project. In terms of time, it took 6.5 years from the discovery of fission to delivering a working bomb. This involved about half a year with no serious government involvement, three years with minor commitment, and three years with major commitment. The term 'Manhattan Project' is sometimes used to refer to these last three years and sometimes to refer to the entire period of US bomb development.[32] I shall use it to refer to the last three years and will say 'US programme' to refer to the entire 6.5 years.

During the 6.5 years, the US programme spent $1.83 billion.[33] This is about $30 billion in 2016 dollars. In relative terms, the peak spending in 1944 was about 0.4% of US GDP, and the total came to about 0.75% of a year's GDP.[34] These costs skewed heavily towards construction, with about two thirds of this money being spent on constructing the facilities and only one third on operating them.[35] The breakdown of these costs are as follows:[36]

---

[31] (Rhodes 1995, p 188).

[32] A major part of the spending in the last three years was under an Army programme code named the 'Manhattan District' or 'Manhattan Engineer District'. The term 'Manhattan Project' derives from these.

[33] Based on the month by month breakdown from (Hewlett & Anderson 1962, p 724) up to the end of August 1945. I have added the $76 million costs of Project Silverplate, which was technically run by the Air Force and not the Manhattan District, but was solely focused on the atomic bomb (the refitting of the bombers, special training of the pilots, and essential logistics).

[34] Most of the cost was borne over the last three years, and particularly in 1944 which alone counted for 51% of the total. I therefore use the 1944 as the base year for inflation and use the 1944 GDP figures.

[35] (Hewlett & Anderson 1962, p 723) and see also (Schwartz 1998). Note that the available cost figures run to the end of December 1945, which unfortunately includes 4.5 months after the end of the war, slightly inflating this estimate of the proportion of operations compared to construction.

[36] (Hewlett & Anderson 1962, p 723). Note that the available cost figures run to the end of December 1945, which unfortunately includes 4.5 months after the end of the war, but it is unlikely the shares were changed too much by the inclusion of this period.

Expense	Share
Oak Ridge (in total)	60%
— K-25 (gaseous diffusion plant)	*26%*
— Y-12 (electromagnetic plant)	*24%*
— S-50 (thermal diffusion plant)	*1%*
— X-10 (test plutonium production plant)	*1%*
— Other	*8%*
Hanford (plutonium production plant)	20%
Special operating materials	5%
Heavy water plants	1%
Los Alamos project	4%
Research & development	4%
Government overhead	2%
Air Force costs (bombers & training)	4%

Table 1. Financial cost breakdown of the Manhattan Project.

This was a lot of money, but it was spent at a time of national emergency and was fairly small compared to the general wartime expenditure. In fact, it was equal to just 9 days of war spending. See the following table for comparisons:[37]

US WWII Expenditure	1940s Dollars	Comparison
Tanks	$5.7 billion	3.1 ×
Artillery	$3.3 billion	1.8 ×
Bombs, mines, and grenades	$2.8 billion	1.5 ×
Small arms (excluding ammunition)	$2.2 billion	1.2 ×
Ammunition for small arms	$2.4 billion	1.3 ×
Clothing	$5.3 billion	2.9 ×
Other	$274 billion	150 ×
Grand Total	**$296 billion**	**162 ×**

Table 2. Partial financial cost breakdown of the US involvement in WWII.

At peak employment (in June 1944), the Manhattan Project was employing a little over 125,000 people — almost one in a thousand people living in the US, and about one in five hundred working people.[38] The total person-years involved come to

[37] (United States Army Services Forces 1946, pp 75–80) and see also (Schwartz 1998). Note that the total war expenditure over the three years was about 150% of a year's GDP.

[38] (United States Army 1948, Appendix 1). The 125,000 figure only includes contractors, notably neglecting several thousand military personnel. Other estimates range from about

around 220,000.[39] Due to bad working conditions, there was a large amount of worker turnover and the total number of people who worked on the project at any point was around 500,000.[40] As with the financial costs, about two thirds of the person years were in construction rather than operation.[41]

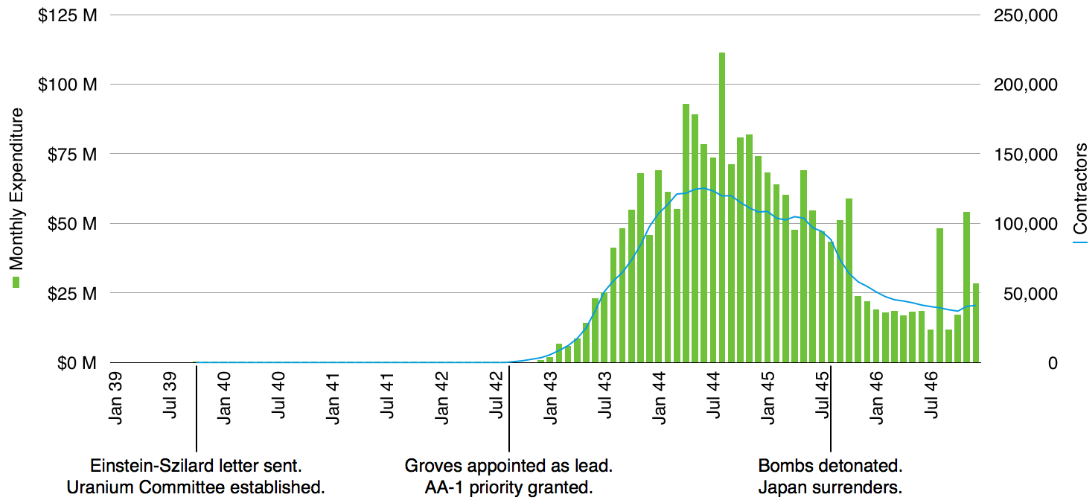

Figure 2. US expenditure and employment from the discovery of fission to the end of 1946.

Figure 2 shows both expenditure and employment trends from the publication of fission in January 1939, to the end of 1946. We can see the three main periods of US government commitment. No commitment until October 1939, when the Einstein-Szilard letter was sent to President Roosevelt and the Uranium Committee established. Then the long period of minor commitment until late 1942. The monthly records for this period are harder to come by, but would barely show up on this graph regardless, as the total expenditure for the period was only $15 million.[42] I have counted the period of major commitment from the appointment of Groves as the project lead and the granting of AA-1 priority, both of which happened in September 1942. This three-year period spent more than 100 times as much as the three years prior. In this report, I have only counted expenditure until August 1945, when the bombs were dropped and the war ended. But I note that the programme did not stop on a dime, and arguably some of the trailing costs should also be counted.

Neither the dollar costs, nor the personnel estimates, capture all the resources involved in achieving all of the pre-requisites for the bomb. One could roughly split pre-requisites into needing 'insight' and 'grind', where the former represents the ability for highly talented individuals to produce sudden progress (such as

125,000 to 130,000. The US population in 1944 was 138.4 million. The US labour force in 1944 was 65.2 million, of which 54.3 was civilian (Long 1952, p 16).

[39] Author's calculation based on chart in (United States Army 1948, Appendix 1).

[40] (Wellerstein 2013). This is about 0.9% of the US workforce at the time.

[41] (United States Army 1948, Appendix 1).

[42] (Hewlett & Anderson 1962, p 724).

theoretical breakthroughs in science or engineering) and the latter represents steady progress achieved with large amounts of money or lower-skilled labour. Obviously there is a continuum between the two, but the labels are useful none-the-less.

Dollar measures for how much each part cost are quite useful, but they underestimate the difficulty of the parts requiring insight. It didn't cost much to hire the scientists — the cost of all the theoretical and practical scientific work on the Manhattan Project was only 8% of the total. However, there was a very limited supply of scientists, especially given that some were needed for other wartime projects such as radar. The high number of great scientists in the US was partly a dividend of having an existing network of well funded universities across a large country, and partly due to the fact that Hitler had forced many Jewish scientists to flee from Europe (which also caused corresponding problems for the German programme). It would have been very difficult for any of the countries to gain many more elite scientists during the six and a half years from the discovery of fission to the end of the war. Scientists were thus a tightly constrained resource, separate from money or lower skill labour.

The Manhattan Project did not perform a careful census of the scientists involved, so unfortunately we have no precise figures for the number of scientists or scientist-years involved. The Atomic Heritage Foundation estimate that over the entire country, between 6,000 and 10,000 scientists were involved, with something like a third to a half being at Los Alamos.[43]

There are also unique resources embodied in individuals who changed the shape of the project. For example, the efforts of Szilard to begin the US political commitment and of Oliphant and Lawrence to reach a major political commitment. Perhaps most important were the exceptional leadership and management abilities of Groves and Oppenheimer.

3. Secrecy

Many scientists realised the potential for bombs as soon as they understood the possibility of the chain reaction, and this moved many of them down the path to secrecy. For instance, by April 1939, the Germans had held a secret meeting on nuclear fission followed immediately by a secret nuclear project in the war office.[44] Similarly, when they realised how small (and thus feasible) the critical mass might be, Frisch and Peierls wrote a secret memorandum to the British government rather than publishing their findings in academic journals.

The person most prominently associated with the push for secrecy was Leo Szilard.[45] He had told few people about the possibility of the chain reaction, and had transferred his patent application to the British Navy to enable it to be kept secret. As

[43] Personal communication from Alexandra Levy of the Atomic Heritage Foundation.

[44] (Rhodes 1986, p 296).

[45] For a detailed analysis of Szilard's work on secrecy, see (Grace 2015).

detailed in Section 1, when he heard the news of fission, he anticipated that others would discover the idea of the chain reaction and was successful in stopping Fermi publishing on the topic, but could not stop Joliot-Curie's team in Paris. As detailed in Section 8, he became a major force for secrecy in the US, successfully stopping Fermi from publishing his work on the feasibility of using graphite as a moderator and stopping Turner from publishing his article about the idea of a plutonium bomb. Eventually the US began systematically preventing the publication of key atomic results, partly due to Szilard's advocacy on the matter.

There was some significant resistance to secrecy among the scientific community. This was mainly from scientists such as Fermi and Joliot-Curie, who wanted to publish their own hard-won discoveries, and thus clearly had a conflict of interest on the matter. Interestingly there was also principled opposition to scientific secrecy, especially from Niels Bohr, who believed it was essential to world progress for science to be completely open and rise above national disputes, seeing secrecy as a grave threat to this vision.[46]

A number of the key discoveries for the building of an atomic bomb ended up being openly published. This began with the discovery of fission, then proceeded with the French papers on the chain reaction, Soviet publication of the moderation properties of graphite and heavy water, then the US (and German) publication about the creation of neptunium. Interestingly, despite being published in prominent journals, the work on moderators and neptunium appear not to have been noticed at the time by the nations that hadn't yet made those discoveries and who would have benefitted greatly from the breach of secrecy. While the Soviets did get some benefit from the US article on neptunium, they only discovered it three years after publication, when their espionage work brought plutonium to their attention.

Even this heightened secrecy itself gave out important information to those who could see it. In 1940, the Soviets realised that the US must be working on a bomb when the names of prominent physicists and chemists started disappearing from the journals.[47] It is not clear whether there was any realistic way for the US to have avoided this unfortunate side effect of their secrecy.

One of the most important secrets was that atomic weapons were feasible at all. While this constitutes just a single bit of information, it was important because there was so much uncertainty about the order of magnitude of resources that would be required and whether the timeline was in years or decades. This naturally made governments reluctant to commit to it, for fear that their effort would be for nothing. In terms of game theory, this was like a 'war of attrition' or an 'all-pay auction', where the government could choose some amount of resources to commit, losing all those resources and if they weren't enough, not getting a new weapon in exchange. Thus the knowledge that it was feasible would allow nations to know that this bet would pay off and that they could scale up their programmes. This was seen in practice when the Soviets scaled up their programme in 1942 and when they scaled it

---

[46] (Rhodes 1986, p 294).

[47] (Rhodes 1986, p 501).

up again immediately after the Hiroshima bombing provided undeniable proof of feasibility.[48]

The Soviet programme was much less secret than the US one — perhaps because they were behind their main rivals, so thought they had much less to lose from spies. Unlike the Americans who hid their main research lab in the desert, the Soviet programme operated out of Moscow. When the Soviet physicist Anatoli Alexandrov first visited the lab, he lost his way and asked a local gang of children for directions. One of them replied: 'It's over the fence where they're making the atomic bomb.'[49]

Knowing that many of the atomic scientists would disperse after the war and that it would be difficult to avoid their knowledge leaking out, Groves commissioned a public report to declassify much of the information. That way, a hard line could be drawn, with key secrets being kept out of the report and scientists who gave away those key secrets being liable for prosecution. The Smyth Report was released on the 12th of August 1945, just three days after the bombing of Nagasaki and three days before the end of the war. It covered many parts of the US bomb development (including isotope separation, reactor building, and making plutonium) but did not mention implosion or certain other key technical details.[50]

This report would have greatly helped many other programmes, but was of relatively little use to the Soviets, as they already knew much of this information. For them it mainly acted as confirmation of the sometimes unreliable espionage reports. There is one particularly interesting exception. In the first release of the Smyth Report, there was a sentence which mentioned the problem of reactor poisoning — one of the key details they intended to leave out. This sentence was thus removed before the second release. However, at some point in the next few months, the Soviets went through the two versions sentence by sentence and were alerted to the importance of reactor poisoning by the discrepancy.[51]

4. Spying

Most of the powers involved in the atomic race did relatively little spying on each others' programmes. The US and the UK were allies and quickly established a partnership, which took away the need to spy on each other. Neither would have gained much by spying on the Soviets or Japanese, but a key motivation for the establishment of both US and UK programmes was the fear of the Germans acquiring unopposed access to nuclear weapons, and so it is somewhat surprising that they appeared to do very little spying on the German programme.[52] After D-

[48] (Rhodes 1986, pp 58–9).

[49] (Rhodes 1995, p 146).

[50] (Rhodes 1995, p 182).

[51] (Rhodes 1995, pp 215–7).

[52] Rhodes provides a good summary of the puzzle (Rhodes 1986, p 605). The fact that they were willing to put so much effort into sabotaging the German heavy water plants makes their absence of spying even more puzzling.

Day, the US did go to great efforts to send spies racing ahead of the troops to find out the state of the German programme, but this was very late in the war.[53] It is difficult to get information about the spying efforts of the Germans or the Japanese, but it appears that little was accomplished.

In contrast, the Soviets excelled at spying, putting many resources into it and gaining an enormous amount of information about how to build an atomic bomb. In 1944, Igor Kurchatov (the head of their atomic programme) noted that the most recent documents he had reviewed totalled about 3,000 pages of text. After the war, another Soviet estimated 10,000 typewritten pages in total.[54] This was not a case of a single snatched document or idea, but meticulous detail on many elements of the weapons programme. While it is impossible to know for sure, it is often estimated that Soviet spying accelerated their progress by one to two years.[55]

By the time of the war, the Soviets had already developed extensive spy networks in other countries. This was greatly helped by the significant interest and optimism about communism in intellectual and working-class circles, with many of the spies being members of communist or Marxist organisations. To get a sense for the scale, at the University of Cambridge (presumably one of their most successful locations), it was estimated that about 150 people were members of communist cells.[56]

While the first Soviet atomic spying was on Germany in 1940 (as they thought them the most likely to develop a bomb),[57] the first big success came from their Cambridge recruits. One of them, John Cairncross, had managed to become private secretary to a minister with access to top secret war files. In September 1941 Cairncross sent the Soviets a summary of the MAUD report.[58] Given the delays in the US, the Soviets thus had this summary of all the British atomic programme's work at about the same time as the American scientists, and perhaps slightly earlier.

While people often view the Soviet spying in the context of the Cold War that followed, it is important to remember that at the time they were allies with the US and the UK. This alliance included the sharing of information. The US and USSR operated a programme called 'Lend-Lease' in which the US sent vast amounts of goods and industrial information to the Soviets to help them build the capacity to hold back Germany. There was thus a somewhat blurry line between sanctioned information for a military ally and espionage material, with the Soviets referring to

---

[53] This programme also involved an earlier phase, looking for evidence within Italy during the Allied invasion.

[54] (Rhodes 1995, p 121).

[55] Rhodes estimates two years (Rhodes 1995, p 74), principally due to finding out about the plutonium pathway in March 1943 rather than in August 1945.

[56] 'Michael Straight, an American student at Cambridge at the time, estimates that "the Socialist Society had two hundred members when I went to Cambridge and six hundred when I left. About one in four of them belonged to Communist cells."' (Rhodes 1995, 51–52).

[57] (Rhodes 1995, p 40).

[58] (Rhodes 1994, pp 52–3). At some later time, the Soviets acquired a complete copy of the report.

the espionage material as 'Super Lend-Lease' and sending it via the same airfield as the regular Lend-Lease material. As a sign of how blurry this line was, the Soviets officially asked for nuclear reactor materials including a kilogram of purified uranium. The US administration granted these requests, though the uranium they sent was deliberately impure.[59] Against this background, we can see that the spies were not helping a national enemy and we can better understand their motivations. Indeed, when the Soviets tried to recruit physicists to give them information on the atomic bomb, they rarely used monetary inducements, but appealed to 'higher feelings' such as the 'good of the world', which they found more effective.[60]

The major pieces of information gained included:

Date	**Spy**	**Information**
Sep 1941	Cairncross	Summary of the MAUD report.[61]
Early 1942	Multiple	The US, UK, and Germany were all working on the bomb.[62]
Mar 1943	Fuchs (?)	Reactors could work without enriched uranium.[63]
Mar 1943	Fuchs (?)	Reactors would make plutonium, which could be used for a bomb.[64]
Feb 1945	Maclean (?)	The US and UK were cornering the world uranium market.[65]
Mar 1945	?	Implosion was a technique that could be useful.[66]
Sep 1945	Greenglass	Information on the implosion lenses and initiator.[67]
Oct 1945	Fuchs	Very detailed information on the implosion bomb.[68]

Table 3. Major pieces of information gained by the Soviets through their atomic spies.

[59] (Rhodes 1995, pp 96–101).

[60] (Rhodes 1995, p 125). Indeed, even attempts to compensate the spies for their time were often met with disgust.

[61] (Rhodes 1995, pp 52–3). This summarised all research by the British to that point in time. In 1943, the Soviets acquired a complete copy of the report.

[62] (Rhodes 1995, pp 58–9).

[63] (Rhodes 1995, pp 70–4).

[64] (Rhodes 1995, pp 74–7).

[65] By the end of the war, the US and UK controlled about 90 per cent of the world's supply of high grade ore in an attempt to deny it to the others. This made things especially difficult for the Soviets, who lacked a good supply, and provided evidence that the Soviet's wartime allies were acting against them. It was probably reported by Donald Maclean (Rhodes 1995, p 130).

[66] (Rhodes 1995, p 152).

[67] (Rhodes 1995, p 187–8).

[68] (Rhodes 1995, p 192–7).

Lessons for new technologies

The historical success of spying may provide a particularly important lesson for the design and development of new transformative technologies. Given how difficult it was to keep the design details secret *even in war time and in top secret committees*, it might be much harder to keep such details secret in peacetime or in a commercial setting. The main objection to this conclusion is that since only the Soviets had such good success, perhaps this was due to very unusual factors that are unlikely to recur.[69] There is definitely something to this objection, but my guess is that spying will continue to be highly effective for those who choose to use it, and very difficult to defend against.

If so, then this would have some particular implications for artificial intelligence. One approach that has been suggested is to focus on making more powerful AI first, and then to sort out the work on safety and control just before it reaches strong general intelligence. However, the case of spying suggests that even if a group increased its security to the level of Los Alamos, key ideas would leak out at a rate similar to the US and UK programmes (in roughly 2 months to 2 years after the breakthrough is made). Rival teams would then be only just behind and there would be almost no lead time to spend on safety and control. Thus, research on control, safety, and ethics, cannot wait for the development of general intelligence, but needs to happen concurrently with direct AI research.

For reference, here are the dates at which key atomic technologies were developed by the first five atomic powers, showing the time lags involved.[70]

	US	**USSR**	**UK**	**France**	**China**
Fission bomb	1945	1949	1952	1960	1964
Fusion bomb	1952	1953[71]	1957	1968	1967
ICBM	1959	1960	1968[72]	1985	1971

Table 4. Dates of development of atomic technologies.

[69] Their heavy investment in a spying infrastructure and the availability of local recruits through communist societies would be good examples. However new ideological commitments such as to openness in technology may play a similar role in promoting leaking of information.

[70] Based on a table from (Bostrom 2014, pp 80–1) which contains several more interesting examples and analysis.

[71] It is unclear what date to use for this. The 1953 test used fusion, but actually got most of its power from fission. A test in 1955 was the Soviets' first 'true' fusion bomb.

[72] Polaris missiles, purchased from the US.

Key Soviet Spies

A lot of the most important information came from a small number of spies. They were thus some of the most important individuals in the development of the atomic bomb.

Klaus Fuchs

- A German physicist who worked on the British Programme from May 1941.
- Recruited by the Soviets while in Britain.
- Sent British atomic secrets.
- Travelled with the British mission to Los Alamos where he stayed from August 1944 to August 1946.
- Passed extremely valuable detailed documents on implosion bomb design and other matters.

David Greenglass

- A US machinist at Los Alamos from August 1944, working on the explosive lenses for the implosion bomb.
- Recruited by the Soviets in November 1944.
- Passed details on the explosive lenses and the initiator to the Soviets in September 1945.[73]

'Perseus'

- An unknown physicist with access to the US's extensive secret scientific literature.
- In spring 1943 he passed 237 abstracts or summaries of important looking atomic research papers, covering isotope separation, reactor design, plutonium, and the chemistry of uranium.[74]

Theodore Hall

- A US physicist who worked at Los Alamos.
- In October 1944 he gave the Soviets a detailed description of the implosion bomb and processes to purify plutonium.

George Koval

- A US citizen recruited by the Soviets by 1939 while living in the Soviet Union.
- He infiltrated the Manhattan Project as a radiation officer in Oak Ridge in 1944.
- Passed details on implosion initiators (and possibly other things) to the Soviets, probably in 1945.

---

[73] (Rhodes 1995, pp 187–8).

[74] (Rhodes 1995, pp 80–1).

Alan Nunn May

- A British physicist based in Canada from 1943 to 1945 where he was involved with the Manhattan Project.
- Recruited by the Soviets while in Canada
- Passed documents and samples of purified uranium-235 and uranium-233.

John Cairncross[75]

- Presumably recruited at Cambridge.
- Become private secretary to a minister with access to top-secret files.
- Supplied the soviets with 5,832 documents (not all atomic).
- Delivered a summary of the MAUD report in September 1941 (the Soviets later acquired a full copy).
- Also worked at GCHQ and sent important decrypts to the Soviets.

Couriers

- Morris & Lona Cohen (couriers for 'Perseus')
- Harry Gold (courier for Fuchs & Greenglass)
- Saville Sax (courier for Fuchs & Hall)
- Julius Rosenberg (courier for Greenglass)
- Ruth Kuczynski (courier for Fuchs in the UK)

5. Estimates

The early thinking about the atomic bomb involved considerable uncertainty about the size of the critical mass and the amount of infrastructure that would be needed to perform a certain amount of isotope separation. Since the cost of a bomb involves multiplying these two numbers, there was extreme uncertainty about whether it would be practical.

As early as May 1939, an estimate of the critical mass of uranium-235 was published in an academic journal by Francis Perrin (an associate of Joliot-Curie in Paris). He estimated a mass of 44 tons as a bare sphere of uranium-235 or 13 tons if surrounded by a neutron reflector. In February 1940, Otto Frisch and Rudolf Peierls were the first to consider the more relevant approach of fast neutron fission in relatively pure uranium-235. They estimated a mass that was 'not a matter of tons, but something like a pound or two.' Further calculations of the size of explosion and cost of production, led them to write the 'Frisch-Peierls Memorandum' which was the first

[75] (Rhodes 1995, pp 52–3).

technical exposition of a practical atomic weapon and led the UK to establish their MAUD committee.

Regarding the amount of infrastructure required, in 1939 Niels Bohr famously said 'It can never be done unless you turn the United States into one huge factory'. While it certainly required a lot of industry (and Bohr joked that he had been right when he saw the scale of the Manhattan Project)[76], its actual peak size of 0.4% of GDP and 0.2% of the national workforce (and 0.02% of the land area) meant that it was realistically achievable.[77]

The MAUD report of July 1941 summarised the UK knowledge to that point in time and offers some useful insight into the quality of the estimates at an early stage (4 years before the bombs were dropped). The estimates of the cost of the bomb turned out to be far too optimistic. They put the cost of a plant capable of producing 10 uranium-235 bombs a year at $20 million.[78] In contrast, the US plants for the uranium-235 bomb cost more than 40 times that amount.[79] The MAUD report put the operating cost per bomb at $1.4 million, but the actual cost to the US was about 70 times higher.[80]

The MAUD report estimated that 10kg of uranium-235 would be needed for a bomb, which would produce an explosion with the energy of about 3.6 thousand tons of TNT. This estimate was better than that of the Peierls-Frisch memorandum,[81] and definitely in the right ballpark, but it is difficult to say much more than that. The Little Boy bomb actually used about 5 or 6 times this much (64 kg of uranium, enriched to 80% uranium-235), but it would have worked with less. Given their use of neutron reflectors, the critical mass for the Little Boy design was less than half as much uranium-235 as they used. Little Boy produced an explosion about 4 times bigger than the MAUD report predicted (~13 kilotons). This estimate was certainly close enough for practical planning purposes, and is very close if interpreted as an estimate of the yield per kilogram or the yield per minimal bomb.

---

[76] (Rhodes 1986, p 500).

[77] See Section 2 for more details and references.

[78] The exchange rate from pounds to dollars was exactly 4:1 for the duration of the war.

[79] It is not trivial to do the breakdown as some plant costs were shared between the uranium and plutonium bombs, but it is at least $870 million (Hewlett & Anderson 1962, p 723).

[80] My calculation comes to $97 million per uranium-235 bomb. This is based on the $291 million operating costs until the end of 1945 (Hewlett & Anderson 1962, p 723), and the fact that at the end of the war, the US was producing enough uranium for a Little Boy bomb every two months (Rhodes 1995, 226–7), giving a rough total of three by the end of the year. The MAUD estimate is less bad if put in terms of kilograms of uranium-235 per dollar rather than bombs per dollar, as the US used 5 or 6 times as much uranium-235 per bomb as the MAUD report predicted.

[81] The memorandum suggested a 5 kilogram bomb, which it predicted would release the energy of 'several thousand tons of dynamite' (dynamite releases about 60% more energy than TNT).

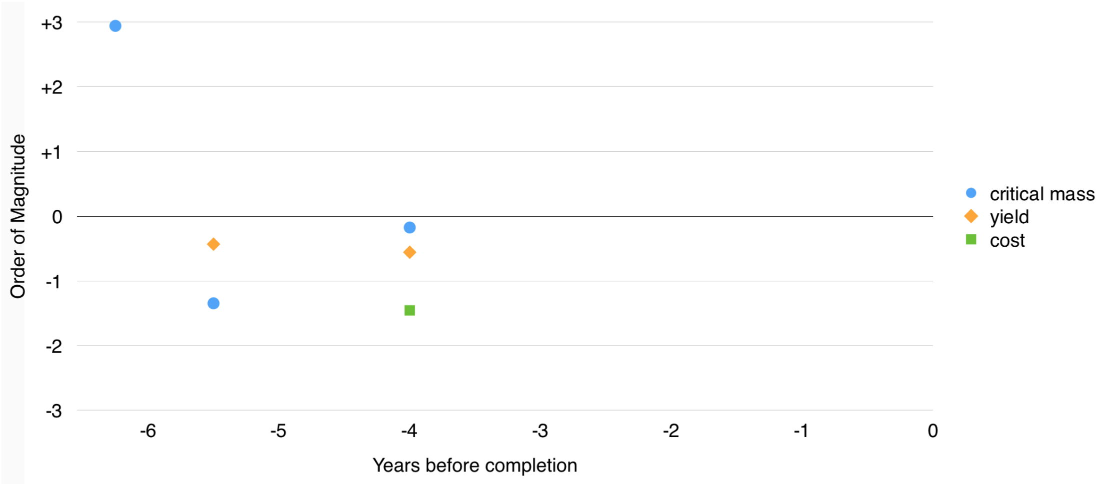

Figure 3. The accuracy of the key estimates as time went on. Zero means perfect accuracy.

From the start of the Manhattan Project, the US was enthusiastically pursuing both the uranium-235 and a plutonium bomb in parallel. As shown in *Figure 1*, these involve very different challenges, with different fissile materials, completely different means of separation, and completely different means of detonation. It is thus highly remarkable that each major branch of the project finished its work almost simultaneously, with the bombs being dropped within just *three days* of each other. If you date the start of serious work on these independent branches to when Groves was put in charge, then this was three days difference at the end of three years of work, or less than 1% difference in time. If you date it to the establishment of the uranium committee, it is even more impressive.

While this may partly be a coincidence, it is also partly the result of very accurate timetabling and prediction. The US wanted a bomb as soon as possible and especially in time to be relevant to the war. They had a very large budget and the ability to put more resources into either type of bomb if it looked like the timetable for that type was slipping and risking irrelevance to the war effort. Doing this correctly involved high accuracy in estimates of the remaining duration of the war as well as in the completion times for all the important parts of each bomb.

6. Deciding how to use the bombs

When Vannevar Bush finally shared the British MAUD report with President Roosevelt, he also explained that in Britain the scientists shared advice on how to use the weapon — something that did not yet happen in the US. Roosevelt decided to explicitly keep it that way, appointing a high-level political committee for nuclear policy. Bush then told the atomic scientists that they were not to consider, discuss, or suggest policy ideas: they had to just help make the bomb.[82]

[82] (Rhodes 1986, p 378). Bush's diaries make clear that one of his aims here was to silence Ernest Lawrence and Arthur Compton.

Trying to lock the scientists out seems to me to have been a large mistake. The development and use of these weapons was to have very serious repercussions for the rest of the century and to threaten the continued existence of the United States (to say nothing of the rest of the world). Additional thought about this at an early stage could have been extremely valuable. Much of the most insightful thinking on atomic weapon policy at the time was from the scientists rather than the politicians — particularly as the latter spent so little time thinking about it. A better compromise would have been to commission some well written advice from them, making it clear that there was no obligation to follow it. An approach like this was used in the Cold War with RAND.

A particularly important set of ideas came from Niels Bohr. Various atomic scientists (most notably Szilard) had suggested that the terrible weapons might usher in an era of peace since any belligerent country would suffer too much to make war worthwhile once the weapons were widespread. Bohr had a clear vision of how a new world order to establish global peace might work, and was very successful at persuading others of his vision. One of his ideas was that the weapon needed to be shared with the Soviets before its use, to bring them in as a partner in a new world order to avert a dangerous arms race. In 1944 he honed a document explaining his views, arranged to deliver it to Roosevelt, and convinced the President of his ideas. Roosevelt sent him to Churchill to try to persuade the British. However, Churchill rejected the idea out of hand as he was against any public concessions by the US or UK, and especially not to the Soviets.[83] In retrospect it is clear that the Cold War was well worth avoiding and that there was little to lose by the plan since the Soviets already had access to the technology through espionage. However it remains unclear how much this would have helped avoid a Cold War, or whether it was a good idea given only the information at hand.

In the months before the use of the bombs in Japan, there was considerable uncertainty and disagreement about their use among both the atomic scientists and the politicians. The main reason to use the bomb was to avoid the greater amount of death that would occur in an all-out invasion of Japan. While not decisive on its own, this did make sense, with many estimates of deaths on each side being far in excess of the number who would die in the bombings. Less noble reasons included to finish the war before the Soviets got involved in order to get a greater share of US control in Japan post-war, and to show their new might to the Soviets in order to secure a better post-war deal in Europe.

In June 1945, the committee that managed the use of the atomic bombs (the Interim Committee) directed its scientific panel to look for some way of convincing the Japanese to surrender via a non-lethal demonstration. The panel, consisting of Lawrence, Fermi, Oppenheimer, and Compton, considered many options suggested by them and their peers, but reluctantly concluded that it could not be done.[84]

[83] (Rhodes 1986, pp 526–38).

[84] (Rhodes 1986, pp 696–7).

In early July 1945, Edward Teller suggested that a bomb used in anger might help people see how bad atomic weapons are and increase the chance of nuclear peace.[85] The prevailing view among the atomic scientists had been that it would set a precedent of using them in anger, which would have had the opposite effect. Presumably both types of effect are real. It is not clear which is stronger.

On 17 July 1945, Szilard and 69 co-signers from the Manhattan Project sent a petition to President Truman advocating for no use of the bomb without giving Japan an opportunity to surrender in light of full information about terms, and for the president to take the effect on starting an atomic arms race into due account when making any decision to drop the bomb. The petition never made it to the President, as it went via the soon-to-be secretary of state, James Byrnes, who was opposed to it. Its main legacy was that it encouraged Groves to take action against Szilard and to fire many of the signatories.

There was also a particularly pernicious reason for dropping the bomb that was appealed to at various points and may have played a role in the decision. Szilard summarised Byrnes' appeal to it in an earlier meeting: 'He said we had spent two billion dollars on developing the bomb, and Congress would want to know what we had got for the money spent.'[86] This is a terrible reason and should have played no part in deliberations. It is either a sunk cost fallacy, vanity, or (at best) the dropping of atomic bombs on civilians to increase one's electoral success.

President Truman ultimately claimed that it was acceptable to drop the bomb given he had duly warned the Japanese and they had failed to surrender. However, unlike the warning Szilard had asked for, Truman's Potsdam Declaration made no mention of the possibility of attack with a novel weapon, so Japan was not in a situation to make an informed decision. Moreover, the Potsdam declaration renewed the Allies' demand for an *unconditional* surrender — something that was not militarily required and which the Japanese were particularly opposed to.

Quite amazingly, the term 'unconditional' only entered into the Allied demands due to a verbal mistake made by Roosevelt when reading a joint statement in a live broadcast in January 1943, a fact that he later admitted.[87] Churchill immediately repeated the demand, later saying: 'Any divergence between us, even by omission, would on such an occasion and at such a time have been damaging or even dangerous to our war effort.'[88] Thus, the otherwise reasonable idea that the bombs needed to be dropped to avoid more deaths in an invasion, was only true due to an unreasonable demand that was created by an error people were too proud to step back from.

[85] (Rhodes 1986, p 697).

[86] (Rhodes 1986, p 638).

[87] '…and the thought popped into my mind that they had called Grant "Old unconditional surrender," and the next thing I knew I had said it.' (Rhodes 1986, p 521).

[88] (Rhodes 1986, p 521).

7. The use of the bombs in Japan

The direct damage caused by the bombs is summarised as follows.

Date	Location	Yield	Casualties
6 Aug 1945	Hiroshima	~13 kilotons	Immediate: 70,000 deaths + 70,000 injured. By five months: 140,000 cumulative deaths. By five years: 200,000 cumulative deaths.
9 Aug 1945	Nagasaki	~21 kilotons	About 60% of the deaths of Hiroshima, due to less suitable location.

Table 5. The direct damage from the atomic bombs.

While these are very high civilian casualties, one should note that they only made up a small proportion of civilian deaths due to bombing in the Second World War. This was because of the vast amount of 'area bombing' with conventional explosives. For comparison, here are the total civilian deaths from area bombing.

Country	Civilian Deaths
Italy	50,000
UK	60,000
China	300,000
Japan	400,000
Germany	400,000
Soviet Union	500,000

Table 6. Total civilian deaths in World War II from conventional area bombing.

Even as individual raids, the Hiroshima and Nagasaki blasts were very large, but arguably not the largest. For comparison:

Date	Location	Duration	Casualties
1940–1	London	9 months	40,000 deaths + 46,000 injured.
Feb 1945	Dresden	3 days	25,000 deaths.
Mar 1945	Tokyo	1 day	100,000 deaths + ~1,000,000 injured.

Table 7. Major area bombing raids in World War II.

So the nuclear bombings were probably the second and third most devastating raids in the war, and Hiroshima may have had the highest eventual death toll.

In terms of raw explosive power per kilogram, the atomic bombs were a discontinuity in the history of explosives.[89] Plutonium is about three million times as explosive as TNT, or two million times better than the best pre-existing explosive. Even if the weight of the rest of the bomb is taken into account, the jump was still dramatic. However, in terms of explosive power per dollar, there was no discontinuity. Supposing we set aside the costs of weapon development and just look at the cost of producing one more Fat Man style bomb, it probably still cost *more* than an equivalent amount of TNT.[90] Perhaps other transformative technologies will follow this pattern: even if there is discontinuous improvement on some metrics, there may just be incremental progress on other equally relevant and important metrics.

The truly remarkable military effect of the atomic bombings was not in terms of raw casualties or casualties per dollar, but the tremendous effect on morale. This was a new and extremely intimidating weapon. It shocked the world.

After hearing news about Hiroshima, Stalin quickly declared war on Japan and his forces attacked Manchuria on 9 Aug 1945, the same day as the Nagasaki bomb.

These developments led the Japanese Emperor to try to arrange a surrender. Despite a failed coup in an attempt to stop the surrender, it was announced 6 days later on 15 Aug 1945. It is unclear whether the Soviets would have invaded were the bomb not dropped and unclear whether the Soviet invasion alone would have been enough to lead the Emperor to surrender.

The next plutonium bomb was ready to be dropped ten days later, and another in ten days after that. Shortly after the war, the estimates were six Little Boy bombs per year and 10 to 12 Fat Man bombs. At that point they decided to stockpile the uranium-235 rather than using it in the obsolete Little Boy design.[91]

8. The plotting of the scientists

A very interesting theme of the US programme was the efforts by some of the scientists involved to shape the course of history. By this I don't just mean their work on critical military technologies, but rather their aims that went beyond science. Starting with Szilard in the early 1930s, they sometimes saw the future possibilities quite clearly and did many things that went beyond standard scientific practice in an attempt to steer things towards a safe and peaceful world.

It is unclear how to assess the impacts of these efforts. The two natural ways are *ex ante* and *ex post* — judging their decisions just in light of the information they had at

[89] See (Grace 2014) for more detail.

[90] (Grace 2014). However, it may have cost somewhat less when including the cost of delivering the explosive to the target, since only a single bomber was required. I don't have any good figures for how much this changes things.

[91] (Rhodes 1995, pp 226–7).

the time and judging them with the full benefit of hindsight. I shall quickly sketch several of the main schemes of the atomic scientists and how they played out.

Early efforts

Szilard had the idea for the chain reaction while living in London in 1933. He very quickly became aware of both the possibility of atomic power and of an atomic bomb (in part due to having read *The World Set Free* by H. G. Wells).[92] He started experiments to find a substance that would undergo a chain reaction, but not knowing about the mechanism of fission, this was very difficult. He was acutely aware of the growing threat of Nazism and was worried about the possibility of a war with Germany, especially if the Germans controlled atomic weapons.[93]

In 1934, he applied for a patent on the idea of fission and when he was later informed that this would need to be assigned to an office of the British government if it was to remain secret, he gave it to the British Admiralty in 1936. Ex post, there were very few strategic consequences of this patent: it played no role in the UK's realisation of the importance of atomic weapons in 1939, and being secret, it did not help the other powers either. Perhaps if things had played out much more slowly, then British ownership of the intellectual property might have been able to shape nuclear developments, but it was such a key technology in a world at war and so secret at first, that intellectual property law turned out to be irrelevant.

Szilard's next major attempt to influence events was when he heard about the discovery of fission while living in New York in 1939. As explained earlier in Section 3, he pushed strongly against the standard scientific approach of sharing all information and made major contributions to increasing the secrecy of the US scientific establishment. At this time he was aided by his fellow Hungarian émigrés Eugene Wigner and Edward Teller, who talked through the issues with him and helped achieve the secrecy. While he notably failed to suppress information about the chain reaction, he did succeed in getting Fermi not to publish his results about the suitability of graphite as a neutron moderator, which (as the Germans had failed to measure this accurately)[94] probably forced the German reliance on their vulnerable heavy water supply.[95]

In the same time period, Szilard also managed to bring the possibility of the atomic bomb to the attention of President Roosevelt (assisted by Wigner, Teller, and Einstein).[96] While this only ended up with a quite limited political commitment, it

---

[92] See Section 11 for more detail on the effects of Wells' writing on Szilard.

[93] (Rhodes 1986, p 223).

[94] (Rhodes 1986, pp 345).

[95] (Rhodes 1986, p 517). Note that it is not entirely clear publication would have helped the Germans, since the Soviets *did* publish their accurate measurements for graphite in a prominent journal and the Germans didn't appear to notice (Rhodes 1995, pp 72–3).

[96] Wigner and Teller pushed very hard for Szilard to reach out to the President (Rhodes 1986, p 303).

did succeed in securing some resources needed for pursuing immediate experiments on the viability of a bomb.

In May 1940, Szilard received a letter from Louis Turner which outlined the theory of how to create what is now known as plutonium for use in a bomb. Turner had prepared a letter to the *Physical Review* and had sent a copy to Szilard for advice on whether it should be kept secret, since Szilard had built a reputation for concern about secrecy. While Turner just saw his method as a way to get explosive power out of the uranium-238 that makes up most of the atoms in natural uranium, Szilard realised that the main asset of the plutonium was that it would be chemically separable from the rest of the uranium, bypassing isotope separation. Szilard wrote back requesting a delay in publication, which Turner agreed to. Presumably if Szilard had not built this reputation and then caught this paper, there is a good chance that Turner's letter would have been published and this would have moved other nations towards the idea of building plutonium bombs.

Racing Germany

A main stated plan of many of the atomic scientists who joined the US and UK programmes was to get atomic weapons before Germany to avoid the terrifying possibility of Germany having unilateral access. For example, the Frisch-Peierls Memorandum (which began the UK programme) stated:

> 'If one works on the assumption that Germany is, or will be, in the possession of this weapon, it must be realized that no shelters are available that would be effective and that could be used on a large scale. The most effective reply would be a counter-threat with a similar bomb. Therefore it seems to us important to start production as soon and as rapidly as possible, even if it is not intended to use the bomb as a means of attack.'

Otto Frisch later clarified:

> 'I have often been asked why I didn't abandon the project there and then, saying nothing to anybody. Why start on a project which, if it was successful, would end with the production of a weapon of unparalleled violence, a weapon of mass destruction, such as the world had never seen? The answer was very simple. We were at war, and the idea was reasonably obvious; very probably some German scientists had had the same idea and were working on it.'[97]

After the war, Einstein gave a similar response, regretting signing the letter to Roosevelt given how things had turned out: 'Had I known that the Germans would not succeed in developing an atomic bomb, I would have done nothing.'[98]

A German atomic bomb was a serious threat given knowledge at the time, and with hindsight too. While the Germans had less industrial power than the US, we do know that the Germans had more political commitment at an early stage (giving them much more time) and that they had sufficient industrial ability (since they

---

[97] (Rhodes 1986, p 325).

[98] 'Einstein, the Man Who Started It All'. *Newsweek*. 10 March 1947.

spent more on their V-2 rocket programme than the US did on the Manhattan Project).

The main thing that might have made it impractical none-the-less was their relative lack of exceptional scientists. This was partly due to the Nazis' anti-Semitic and anti-intellectual policies driving many of them from the country (something that Max Planck had personally warned Hitler about in 1933). For example, 10 physicists and 4 chemists who had won or would win Nobel prizes emigrated from Germany shortly after Hitler came to power.[99] Half the German nuclear physicists who were cited in the literature before 1933 emigrated.[100] As well as disadvantaging the German program, these scientists in exile were part of what helped the US and UK programmes excel. For instance, both authors of the Peierls-Frisch memorandum (which catapulted the UK programme into action) were Germans in exile.

Among the achievements of the German programme were understanding the potential use of plutonium for a bomb by July 1940 and demonstrating positive neutron production in a reactor by April 1942 (three months before the US).[101] After the war, when he saw how the US had scaled up its programme, Heisenberg realised that the achievement of positive neutron production had been a pivotal moment for the German programme: it was proof that a self sustaining reaction was within reach and therefore that a case could be made to really scale up their programme. However, the physicists did not make this case, and due to an accident in their prototype reactor and the deprioritization of their research, the German programme did not even manage to achieve criticality.[102]

While we now know that the Germans didn't get close to a working bomb, racing to build a bomb before them does seem to have been a reasonable way of making sure Hitler didn't have unrivalled access to atomic weapons. It did have somewhat predictable downsides (with the signs that the US was working on a bomb causing some minor acceleration of the German programme) and so it probably increased the chance that the Germans would develop a bomb. But it wasn't unreasonable to be most worried about unilateral German access to a bomb, which their work did combat.

Later efforts

As the bomb development continued, many atomic scientists did start to worry about the longer term effects of causing an arms race with the Soviets that could threaten continued human civilisation itself. The next really notable attempt by atomic scientists to change the course of history was Bohr's idea for bringing the Soviets on as full partners in the hope of avoiding an arms race (see Section 6). Among scientists Bohr was second only to Einstein in fame and was an excellent statesman, managing to sell the idea to Roosevelt, before it was ultimately blocked

[99] (Beyerchen 1977, p 48).

[100] (Hentschel & Hentschel, 1996, p lviii).

[101] (Kant 2002, pp 7–8).

[102] (Carson 2010, p 372).

by Churchill. While we will never know whether Bohr's plan would have avoided a cold war, with the benefit of hindsight we can see that it had relatively little downside since the Soviets already had access to the cutting edge US atomic secrets through espionage. Given the failure of Bohr's plan, and the successful Soviet spying, we can now see that the technical work of the atomic scientists ended up doing quite a lot to accelerate Soviet atomic weapons development and to plunge the world very quickly into a cold war.

Section 6 detailed the last minute efforts by the atomic scientists in 1945 to find ways of avoiding the use of the bombs. These included the attempt of a panel of scientists to find a suitable non-lethal demonstration and of Szilard's petition for the US to give Japan an informed opportunity to surrender and to fully consider the chance of starting an arms race when weighing whether or not to use the bomb. Both of these were well thought out attempts, but ended up failing.

Finally, a major legacy of the strategic thought of the scientists involved is the Bulletin of the Atomic Scientists, an influential journal aimed at informing the public and policy makers about the threat of nuclear war. More recently, it has added other technological threats to its purview and is a leading scientific journal for thought on man-made existential risks.

In summary, the atomic scientists had significant foresight about the major issues, and tried very hard to influence the world for the better, achieving some notable successes, but not obviously creating good effects overall.

9. The effects of individuals

We know that the development and deployment of key technologies can have a transformative impact upon the world, but it is less clear what role individuals can have in the process. For example, what impact could an individual realistically have on reducing the risk of a catastrophic outcome during a major technological transition (such as the development of atomic weapons, or of advanced artificial intelligence)?

The extremely well documented history of the development of the atomic bomb gives us some insight into this. There were many individuals who appeared to have a dramatic effect on the development, potentially changing the timeline for a working bomb by up to a year or more, changing the chance of programmes reaching success at all, and changing how the bomb would be deployed. While we can never know for sure what would have happened if individuals had acted differently, the details of the cases I outline below are suggestive, and in some cases we know that things played out differently in the other nations' programmes.

While it is very difficult to convey this evidence in summary, I'd strongly recommend reading a history of the period (such as Rhodes' *The Making of the Atomic Bomb*) to get a feel for how sensitive the process was to the actions of individuals.

Some of the most influential people (in roughly chronological order):

Leo Szilard *(and his conspirators Eugene Wigner and Edward Teller)*

- Played the largest role in getting US government involvement.
- Successfully pushed for more secrecy about key results (including graphite and the plutonium bomb).

Enrico Fermi

- Led the work on reactor design, making very large contributions.
- Fought against secrecy during the crucial early months.

Edwin McMillan

- Discovered neptunium (which decayed to plutonium).
- Published this, giving away a major atomic secret (though the effects of this were smaller than feared).

Lyman Briggs

- Led the uranium committee that delivered extremely slow progress for the first three years of the US programme.
- Kept the UK's extremely important MAUD report away from the US scientists until he was forced to share it five months later.[103]

Vannevar Bush

- Along with Briggs, he was also very influential in blocking acceleration of the US atomic programme.

Mark Oliphant

- Flew from the UK to the US to find out why they were ignoring the MAUD report and spurred them to action.[104]

Ernest Lawrence

- Developed the cyclotron and adapted it for electromagnetic separation.
- Successfully pushed for acceleration of the US programme

[103] (Rhodes 1986, p 372).

[104] (Rhodes 1986, p 372). On the same page, Szilard is quoted after the war as saying: 'If Congress knew the true history of the atomic energy project, I have no doubt but that it would create a special medal to be given to meddling foreigners for distinguished services, and Dr Oliphant would be the first to receive one.'

Leslie Groves

- An amazingly quick and effective administrator for the US programme. His appointment was a major turning point.

Robert Oppenheimer

- Expertly led the research programme at Los Alamos.

George Kistiakowsky

- Helped the US confirm scientific importance of nuclear research early on.
- Spearheaded the essential work on the implosion bomb.

John Cairncross

- Passed a summary of the MAUD report to the Soviets.

Klaus Fuchs

- Passed many secrets to the Soviets, including detailed knowledge of implosion.

Niels Bohr

- Devised a promising method for avoiding a cold war by bringing the Soviets on as partners in developing the bomb.

Franklin Roosevelt

- Saw the potential of the bomb and attempted to expedite it.
- Accidentally called for *unconditional* surrender of the Japanese, leading to the eventual need for the bomb to be dropped.
- Saw the potential of Bohr's plan and sent Bohr to talk to Churchill about it.

Winston Churchill

- Quickly dismissed Bohr's plan due to distrust of the Soviets.
- Repeated Roosevelt's call for unconditional Japanese surrender, making it much more difficult to back down from the mistake.

Harry Truman

- Decided to use the bomb on Japan (without any real warning).

Note that the above list is, by necessity, rather subjective and presumably biased towards those who actions were more easily observed. It is also somewhat biased towards whichever people turned out to have large effects, rather than directly reflecting how much influence a concerned and dedicated actor could expect to have (though Section 8 explored that in more detail). Regardless, this list should still provide a useful idea of the contingency of historical events upon individual behaviour and a list of people whose roles may warrant deeper investigation.

10. Existential Risk

It is well known that the atomic bomb may be the first technology that could destroy the human race itself. An all out nuclear war between nations with large stockpiles of advanced nuclear weapons would lead to the burning of many large cities and the creation a vast amount of soot in the upper atmosphere. This may lead to global cooling and crop failure. While we do not know exactly how bad this effect would be or how long this would last, there is a live possibility that the nuclear winter might lead to human extinction.

It is less well known that there were serious concerns that the mere development of the bomb might lead to the accidental destruction of humanity.[105] In the summer of 1942, a group of US scientists led by Oppenheimer and Teller investigated the possibility of developing a bomb where an initial fission explosion set off a thermonuclear explosion in an attached payload of deuterium (which we now know as the hydrogen bomb or fusion bomb). While doing so, Teller noticed that the explosion of a fission or fusion bomb would create a heat that was unprecedented upon the earth. This created a theoretical possibility that the heat might allow an additional self-sustaining thermonuclear reaction in the surrounding air or water, which could continue to spread, enveloping the entire world in flame. Thus it might not take extreme warlike actions from major nations to cause destruction, but a single nuclear test during bomb development, before the project had even undergone any public scrutiny. Moreover, if Teller were right, then this disaster could be triggered by a secret German programme just as well as by the US programme.

While Teller's concerns were not taken seriously by the entire group, Oppenheimer recognised their importance and informed his superior Arthur Compton who recollected:

> 'Was there really any chance that an atomic bomb would trigger the explosion of the nitrogen in the atmosphere or the hydrogen in the ocean? This would be the ultimate catastrophe. Better to accept the slavery of the Nazis than to run a chance of drawing the final curtain on mankind!'[106]

Further calculations were undertaken immediately, which convinced Teller and the others that a runaway explosion would not in fact be possible, and that it only appeared possible due to very conservative assumptions that Teller was using.

From an existential risk perspective, it is great that the team took this concern seriously and put considerable effort by top scientists into checking whether it was physically possible. No doubt some scientific teams would have simply dismissed the concerns. However, it is unfortunate that this checking took place in poor

---

[105] See (Rhodes 1986, pp 418–9).

[106] (Rhodes 1986, p 419).

epistemic conditions. Due to wartime secrecy the calculations were not subject to any external review,[107] and since atomic weapons would be of great benefit to their country the scientists were vulnerable to motivated thinking. Arguably, it was not for a single nation (let alone a handful of scientists) to decide for the rest of humanity whether or not their checking had been thorough enough.

Lest we be too confident in the calculations of top atomic scientists, it is worth considering the Castle Bravo test of 1954.[108] This was one of the early thermonuclear tests and used lithium as its main fuel. Like uranium, this came in two main isotopes. When calculating the expected yield of 5 megatons (350 times that of Hiroshima), the scientists at Los Alamos assumed that the lithium-6 would contribute while the lithium-7 would not. They thus removed much of the lithium-7 with isotope separation before the test. However it turned out that their calculations were in serious error, as the remaining lithium-7 contributed more energy than the lithium-6. The bomb exploded with three times the predicted yield (more than 1,000 times that of Hiroshima), greatly increasing the irradiated area and causing a minor international incident. The scale of the explosion was never exceeded by subsequent US nuclear tests and it is known as the highest energy mistake in human history. Indeed, the supposedly inert lithium-7 became the standard main ingredient in subsequent US thermonuclear bombs.

This mistake was on a calculation the US atomic scientists had spent much longer developing, and in a situation where they had more knowledge of thermonuclear reactions. I am thus not too sure that we can say the US scientists successfully managed the potential existential risk that Teller noticed, since we should demand that the chance of a major calculation error was less than 1 per cent, yet such an error later happened in a somewhat similar situation.

While it was impossible to completely overcome these issues, there were approaches that could have been taken. For instance, a 'red team' could have been created and tasked with coming up with the strongest arguments for not proceeding with a bomb. These could include arguments that the sustained reaction would have been possible, or why it *might* still be possible, or why even a small remaining chance might be too great a risk. Additionally, to avoid group think, a separate team of atomic scientists could have been tasked with analysing the problem, to see if they independently came up with the same reasons for its safety. In the absence of proper peer review or international decision making about whether the risk had been adequately assessed, these approaches might have at least improved the robustness of the decision making.

Even fewer people are aware that the Germans scientists also considered the possibility that atomic weapons might cause our extinction through an uncontrolled continuing chain reaction. In their case, it was fission they were worried about. Albert Speer, the German minister of armaments later gave a chilling account:

---

[107] Thus the typical academic paper is subject to more external review than this potentially world ending technology received.

[108] My colleagues and I discuss this situation in some detail in (Ord et al 2010).

> 'Hitler did sometimes comment on [the prospects of nuclear fission], but what I told him of my conferences with the physicists confirmed his view that there was not much profit in the matter. Actually, Professor Heisenberg had not given any final answer to my question whether a successful nuclear fission could be kept under control with absolute certainty or might continue as a chain reaction. Hitler was plainly not delighted with the possibility that the earth under his rule might be transformed into a glowing star. Occasionally, however, he joked that the scientists in their unworldly urge to lay bare all the secrets under heaven might some day set the globe on fire. But undoubtedly a good deal of time would pass before that came about, Hitler said; he would certainly not live to see it.'[109]

11. Additional lessons and insights

The following are an assortment of smaller lessons and insights from the development of atomic weapons.

The structure of the challenges

It is striking that there were several very different ways of proceeding which were all ultimately viable for producing a bomb (corresponding to the branches in *Figure 1*). This includes both the large split between the two types of bomb, and then the variety of methods for producing the fissile material for each type. Many programmes only knew about a few of these routes, which limited them to suboptimal options. It was thus particularly valuable to find out about additional routes.

Pursuing all branches

The successful US programme knew about all the branches in *Figure 1* and took a very interesting approach to them. Rather than focusing resources on the most promising branch (keeping the others as fall-backs) they repeatedly pursued all branches in parallel.[110] This was possible because of their relatively unconstrained budget and because the long processing times in the scientifically easier project (the uranium-235 bomb) meant that scientific talent wasn't always needed there, so could be rather cheaply redirected to the plutonium bomb. The US knew that the timing constraint of developing a bomb (before the Germans and before the end of the war by other means) was tighter than the financial constraints. Their timetabling was so successful that the two very different weapons with their very different challenges were ready for deployment within just three days of each other.

[109] (Rhodes 1986, pp 404–5).

[110] Indeed, one of their only mistakes appears to have been prematurely giving up on centrifuge isotope separation (Kemp 2012).

Re-checking key assumptions

Figure 1 shows the structure of the pre-requisites, complete with alternative paths. It is striking that in several cases mistaken assumptions or measurements led to neglecting important paths, or pursuing them in ways that couldn't be made to work. For example, the US could have pursued fast neutron or slow neutron uranium-235 bombs, but appears to have assumed slow neutron bombs would be superior without ever having really questioned this assumption (they were eventually told their mistake by the UK programme in 1941). The Soviets assumed that a uniform mixture of uranium and its moderator would be optimal for reactor design and only discovered this crucial mistake through espionage in 1943.[111] The UK and German programmes made mistakes in the relatively easy process of measuring the moderation properties of graphite, which caused them to write off its use as a moderator. The UK was told of their mistake by the US, while the Germans never found out, doubling down on heavy water and making their programme vulnerable to the repeated Allied sabotage of their only heavy water supply.

It also appears that the failure of the US and USSR to develop efficient centrifuges for isotope separation was not because the required technology was too difficult, but because they went down design paths that would not work, and kept making iterative changes to them rather than rethinking the main element of the design.[112]

These were all cases where the key steps were not very difficult in objective terms, but failure to get them right the first time, and then never re-checking them, was a major setback.

Alliance

The effective merger between the UK and US programmes was extremely helpful for the allied efforts (both the sharing of the MAUD report and then the sending of scientists to Los Alamos in 1943). While the UK did not have the industrial capability during the wartime to produce the fissile material for a bomb, they were at the forefront of the science and made substantial contributions to the scientific and engineering basis of the US programme. They were also responsible for significantly increasing the scale of the US political commitment. It seems very unlikely to me that the US would have been able to produce a bomb during the war without the help of the UK programme.[113] Similar mergers may be very helpful if there are races towards other transformative technologies.

[111] The Germans made a similar mistake, assuming that layered arrangements were best.

[112] (Kemp 2012).

[113] Groves has given statements on this, but they appear to contradict each other. He has said British assistance was 'helpful but not vital' but also that "without active and continuing British interest, there probably would have been no atomic bomb to drop on Hiroshima.' (Groves 1962, p 408).

However, note that while helping succeed in creating the bomb, the UK scientists in Los Alamos included Klaus Fuchs, who was then able to give the secrets of implosion to the Soviets, making the post-war period much harder for the US.

'Conservative assumptions'

At the early stage, there was great uncertainty about whether a bomb would be feasible. Given such uncertainty, scientists and engineers are often drawn to the idea of making 'conservative' assumptions as a way of exploring the uncertain situation via a realistic worst case. However, there are actually two types of conservatism which are in conflict with each other. In typical scientific situations, being conservative means assuming your results will not overturn the scientific paradigm nor have dramatic results on the world stage. In situations where a great deal is at stake, by contrast, being conservative means assuming things might be more dangerous than your central estimate implies. This ambiguity came up when Szilard was trying to convince Fermi to keep the idea of the chain reaction secret:

> 'From the very beginning the line was drawn … Fermi thought that the conservative thing was to play down the possibility that [a chain reaction] may happen, and I thought the conservative thing was to assume that it would happen and take all the necessary precautions.'[114]

The clash between these two types of conservative assumption has also come up in recent discussions of the safety of AI (such as the 2015 Puerto Rico conference), with practitioners of AI naturally making the 'conservative' assumption that they will not make progress more rapidly than expected, and safety experts making the 'conservative' assumption that they will. This causes misunderstandings when the assumption is in the background and communication difficulties when they are both using the word to mean different things.

Unilateral action

Some actions, such as deciding to publish potentially dangerous research or to deploy a risky technology in the field are subject to a problem that has been called the *unilateralist's curse*.[115] In such situations where it is possible for any party to act unilaterally, there is a problem where the different parties will have different estimates of whether the benefits will outweigh the costs. Even if most of them agree that it would be bad to proceed, it is the most optimistic estimate that ends up determining whether the action is taken. There is thus a bias towards taking action. Bostrom et al argue that a useful practical rule in such situations is to pool estimates and take the median one (or equivalently, to vote on the matter). This is shown to perform almost as well as the optimal Bayesian solution.

Interestingly, exactly such a method was used in New York when Szilard, Teller, and Fermi met in March 1939 to discuss whether to tell the world about the chain reaction. As Szilard later told it, they met:

---

[114] (Rhodes 1986, p 281).

[115] By analogy to the *winner's curse* in auctions and tendering (Bostrom et al 2016).

> 'to discuss whether or not these things [the *Physical Review* papers on the possibility of a chain reaction] should be published. Both Teller and I thought that they should not. Fermi thought that they should. But after a long discussion, Fermi took the position that after all this was a democracy; if the majority was against publication, he would abide by the wish of the majority.'[116]

In this way, Fermi avoided falling victim to the unilateralist curse.

Science fiction quickly becoming science

Like several other key technological developments, the atomic bomb was foreshadowed in science fiction. In 1932, a year before he discovered the chain reaction, Szilard read H. G. Wells' *The World Set Free*.[117] This prescient 1913 work covered many of the key ingredients that would shape the strategic situation to follow. It included weapons based on releasing atomic energy, that gave such an advantage to the aggressor's initial assault that they necessitated the formation of a world government to control their use, and thus brought about a world peace. Having already encountered these ideas must have helped guide Szilard's thoughts when a very similar scenario started to play out around him.[118] As he would later write: 'Knowing what [a chain reaction] would mean — and I knew because I had read H.G. Wells — I did not want this patent to become public.'[119]

Immediately prior to Szilard's idea of the chain reaction, the scientific establishment was still dismissing the idea of gaining useful amounts of atomic power as 'nonsense'. Indeed, just a day before Szilard discovered the chain reaction, Ernest Rutherford — pre-eminent atomic scientist of the age — gave a lecture where he said 'Any one who expects a source of power from the transformation of these atoms is talking moonshine.'[120] As Szilard later described:[121]

> 'In the fall of 1933, I found myself in London. I kept myself busy trying to find positions for German colleagues who lost their university positions with the advent of the Nazi regime. One morning I read in the newspaper about the annual meeting of the British Association where Lord Rutherford was reported to have said that whoever talks about the liberation of atomic energy on an industrial scale is talking moonshine. Pronouncements of experts to the effect that something cannot be done have always irritated me. That day, I was walking down Southampton Row and was stopped for a traffic light. I was pondering whether Lord Rutherford might not prove to be wrong. As the light changed to green and I crossed the street it suddenly occurred to me that if we

---

[116] (Rhodes 1986, p 295).

[117] It had also been read by Churchill (who met with Wells on several occasions) and may have also influenced his thinking about atomic matters.

[118] Szilard later wrote that while *The World Set Free* had a large effect on him, it didn't cause him to immediately go into nuclear physics to pursue these ideas (Rhodes 1986, p 24). It was after he made this theoretical discovery that it had its influence.

[119] (Szilard 1978).

[120] (Rhodes 1986, p 27).

[121] (Szilard 1972, p 529).

> could find an element which is split by neutrons and which would emit two neutrons when it absorbed one neutron, such an element if assembled in sufficiently large mass, could sustain a nuclear chain reaction, liberate energy on an industrial scale, and construct atomic bombs. The thought that this might be possible became an obsession with me. It led me to go into nuclear physics, a field in which I had not worked before.'

There is some room for dispute about the exact timing of Szilard's revelation (whether it happened that very day or a few days later) and in exactly how clear the chain reaction was to him at the time (for example whether he had thought of the exponential chain reaction at the time rather than a merely linear one).[122] Moreover, the idea of the chain reaction was still a long way from a practical implementation of it (9 years). But the basic fact still stands that this was a quick transition from disreputable science fiction to transformative technology — especially as Rutherford repeated the 'moonshine' claims until his death in 1937.[123]

As the history of technology is littered with examples of ideas that were publicly ridiculed by the most prominent scientists shortly before they became reality, there is a very real prospect this will also be the case in the lead up to new transformative technologies. Of course this doesn't mean that all ideas from science fiction will be implemented, but it does mean that strong dismissive claims by prominent experts provide much less evidence against them than you might think.

Racing

A serious concern regarding the development of transformative technologies is that a dangerous race might arise between competing teams. This is most well known in cases where the technology has military applications, leading to arms races. However, it could also be dangerous in cases with large non-military benefits if the technology's development or deployment has potential global risks. It has been suggested that this might be the case for artificial intelligence: if there is a winner-take-all dynamic, then each team has a strong incentive to get there first and would be tempted to cut corners on safety if that helped them beat their rivals.[124] It was also suggested that sharing information between the competing groups might make the risk race even worse.[125]

Interestingly, there was an episode in the making of the atomic bomb that showed the opposite happening. In late 1942, after achieving criticality in a reactor, the US

---

[122] For a very good analysis, see (Wellerstein 2014).

[123] In Rutherford's defence, it has been suggested that he did see the possibility of atomic power and wanted to dissuade people from attempting to harness it, in order to prevent atomic weapons. There is some evidence to back this up, including that he expressly warned a minister in the War Office about the use of atomic weapons, saying he 'had a strong hunch that nuclear energy might one day have a decisive effect on war' (Jenkin 2011). However, whether or not Rutherford privately worried about this does not change the fact that the public pronouncements of the scientific establishment were of poor guidance.

[124] (Armstrong et al 2013).

[125] (Armstrong et al 2013).

programme began to think about the possibility of using highly radioactive materials from a reactor to poison one's enemies.[126] Since they had very little information about the German programme, they (correctly) feared that it might have got started more quickly and (incorrectly) feared that it might have achieved criticality years earlier. Even if the Germans hadn't yet mastered a bomb, they might have been able to store enough radioactive material to seriously poison allied food supplies. By mid 1943, some of the foremost atomic scientists were sufficiently scared that they began the theoretical work on developing such weapons themselves. They went so far as to determine which isotopes has the best combination of ease of production, high radioactivity, and uptake into a person's organs, as well as how best to distribute the poison. Happily, further analysis suggested that radioactive poisoning would be an unlikely approach for the Germans to take and US research on it ceased. However, it does show how lack of information about one's rivals can also accelerate dangerous developments.

One could also consider the rivalry between the US and the Germans more broadly. Since many US scientists were mainly motivated by the fear that the Germans might develop the bomb, it seems that some scientists might have stopped working on the bomb at all if they found out that the German bomb programme had been discontinued in 1942. However, this wouldn't count very strongly against the general idea that information sharing can make things worse, for more information about the Germans *prior* to 1942 would probably have accelerated US developments, and if the *Germans* had had good information about the US programme in 1942, they would very likely have scaled up their bomb development instead of discontinuing it.

Overall, news that the rival teams are close would typically accelerate things and news that they are not would decelerate them. Which effect dominates on average is a difficult question, depending both on how likely it is that they are close, and on the psychological effects of fear of the unknown.

Beyond 1945

To keep things manageable, I have mostly limited my analysis of the development of atomic weapons to the period up to the end of the Second World War. However, one could continue this analysis through the Cold War. In my view there is enough material there for another report of about this length. This could include coverage of:

- Continued technological developments (such as the fusion bomb, nuclear submarines, centrifuges, ICBMs, and MIRV).
- Continued espionage.
- The Baruch plan.
- Attempts to avoid proliferation.
- Nuclear incidents where weapons were almost deployed (including accidents).

[126] (Rhodes 1986, pp 510–1).

- The role of the atomic scientists.
- The role of game theory and political science.
- The role of RAND.

Other relevant case studies

In this report, I have explored the development of atomic weapons as an important case study for thinking about the development of other transformative technologies. It is my belief that it is among the most informative examples to study. However care must be taken not to over fit to this one case. Things will not play out exactly this way again, and one would get a better picture of the space of possibilities by studying several different cases.

Other cases with clearly defined teams working urgently towards important (though perhaps not world-shaping) technologies include:

- The space race
- The code breaking at Bletchley Park
- Radar
- Human genome sequencing

Alternatively, the development of new transformative technologies might proceed in a significantly less dramatic manner. For example, it might be something with much less competitiveness such as the green revolution, or something with much more dispersed and incremental progress such as semiconductor development. While there is much to learn from the history of atomic weapons development, a careful analyst should look into other cases too.

References

AI Impacts. 2014. Discontinuity from Nuclear Weapons, AI Impacts, accessed March 19, 2016, http://aiimpacts.org/discontinuity-from-nuclear-weapons/.

Armstrong, Stuart, Nick Bostrom, and Carl Shulman. 2013. Racing to the precipice: a model of artificial intelligence development, Technical Report #2013-1, Future of Humanity Institute, Oxford University: pp. 1-8.

Beyerchen, Alan D. 1977. *Scientists Under Hitler: Politics and the Physics Community in the Third Reich*. New Haven: Yale University Press.

Bostrom, Nick. 2014. *Superintelligence*. Oxford: Oxford University Press.

Bostrom, Nick, Thomas Douglas & Anders Sandberg. 2016. The Unilateralist's Curse and the Case for a Principle of Conformity, Social Epistemology, DOI: 10.1080/02691728.2015.1108373

Carson, Cathryn. 2010. *Heisenberg in the Atomic Age: Science and the Public Sphere*. New York: Cambridge University Press.

Grace, Katja. 2015. *Leó Szilárd and the Danger of Nuclear Weapons: A Case Study in Risk Mitigation*. Technical report 2015–10. Berkeley, CA: Machine Intelligence Research Institute.

Groves, Leslie. 1962. *Now it Can be Told: The Story of the Manhattan Project*. New York: Harper & Row.

Hentschel, Klaus and Ann M. Hentschel. 1996. *Physics and National Socialism: An Anthology of Primary Sources*. Basel: Birkhäuser.

Hewlett, Richard G., and Oscar E. Anderson Jr. 1962. *The New World: A history of the United States Atomic Energy Commission*, Volume 1, 1939/1946. University Park, PA: The Pennsylvania State University Press.

Holloway, David. 1994. *Stalin and the Bomb: The Soviet Union and Atomic Energy, 1939–1956*. New Haven, CT: Yale University Press.

Long, Clarence D. 1952. The Labor Force in Mobilization: United States, Great Britain, Germany, in Clarence D. Long (ed.) *The Labor Force in War and Transition: Four Countries* NBER, pp. 8–32. http://www.nber.org/chapters/c3132

Jenkin, John G. 2011. Atomic Energy is "Moonshine": What did Rutherford Really Mean?. *Physical Perspectives* vol. 13 (2) pp. 128-145.

Kant, Horst. 2002. Werner Heisenberg and the German uranium project. Preprint 203. Berlin: Max Planck Institute for the History of Science.

Kemp, R. Scott. 2012. The End of Manhattan: How the Gas Centrifuge Changed the Quest for Nuclear Weapons. *Technology and Culture* 53(2), 272–305.

Ordway, Frederick, and Mitchell R Sharpe. 1982. *The Rocket Team*. Apogee Books Space Series 36.

Rhodes, Richard. 1986. *The Making of the Atomic Bomb*. London: Simon & Schuster. Page numbers are from the 2012 paperback edition.

Rhodes, Richard. 1995. *Dark Sun: the making of the hydrogen bomb*. London: Simon & Schuster. Page numbers are from the 2005 paperback edition.

Schwartz, Stephen I. 1998. *Atomic audit: the costs and consequences of U.S. nuclear weapons since 1940.* Washington, DC: Brookings Institution Press.

Szilárd, Leó. 1934. Improvements in or relating to the transmutation of chemical elements GB630726, filed 1934. Accessed 16 Feb 2016. http://worldwide.espacenet.com/publicationDetails/biblio?CC=GB&NR=630726.

Szilard, Leo. 1972. *The Collected Works: Scientific Papers.* Cambridge, MA: MIT Press.

Szilard, Leo. 1978. *Leo Szilard: His Version of the Facts. Selected Recollections and Correspondence.* Cambridge, MA: MIT Press.

Ord, Toby, Rafaella Hillerbrand, and Anders Sandberg. 2010. Probing the improbable: methodological challenges for risks with low probabilities and high stakes. *Journal of Risk Research* 13 191–205.

United States Army, Corps of Engineers, Manhattan District. 1948. *Manhattan District History*. Volume 8 — Personnel. Washington, DC: US Department of Energy, Office of History and Heritage Resources.

United States Army Service Forces. Control Division. 1946. *World War II: A Summary of ASF Statistics.* Washington, DC: Statistics Branch, Control Division, Headquarters, Army Services Forces, War Department.

Wellerstein, Alex. 2013. How many people worked on the Manhattan Project?, Restricted Data: The Nuclear Secrecy Blog, accessed January 25, 2016, http://blog.nuclearsecrecy.com/2013/11/01/many-people-worked-manhattan-project/.

Wellerstein, Alex. 2014. Szilard's chain reaction: visionary or crank?, Restricted Data: The Nuclear Secrecy Blog, accessed March 10, 2016, http://blog.nuclearsecrecy.com/2014/05/16/szilards-chain-reaction/.

Centre for the
Governance of AI