JOURNAL OF INVESTMENT MANAGEMENT, Vol. 21, No. 1, (2023), pp. 5–51

www.joim.com

FINANCING FUSION ENERGY

Abdullah Alhamdan[a], *Zachery M. Halem*[b], *Irene Hernandez*[c], *Andrew W. Lo*[d,*], *Manish Singh*[e] *and Dennis Whyte*[f]

The case for investing in fusion energy has never been greater, given increasing global energy demand, high annual carbon dioxide output, and technological limitations for wind and solar power. Nevertheless, financing for fusion companies through traditional means has proven challenging. While fusion startups have an unparalleled upside, their high upfront costs, lengthy delay in payoff, and high risk of commercial failure have historically restricted funding interest to a niche set of investors. Drawing on insights from investor interviews and case studies of public–private partnerships, we propose a megafund structure in which a large number of projects are securitized into a single holding company funded through various debt and equity tranches, with first loss capital guarantees from governments and philanthropic partners. The megafund exploits many of the core properties of the fusion industry: the diversity of approaches to engender fusion reactions, the ability to create revenue-generating divestitures in related fields, and the breadth of auxiliary technologies needed to support a functioning power plant. The model expands the pool of available capital by creating tranches with different risk–return tradeoffs and providing a diversified "fusion index" that can be viewed as a long hedge against fossil fuels. Simulations of a fusion megafund demonstrate positive returns on equity (ROE) and low default rates for the capital raised using debt.

1 Introduction

With greenhouse gas levels rising, the effects of global warming already tangible in the form of extreme weather events, and the world's supply of fossil fuels expected to be depleted by the

[a]MIT Sloan School of Management
[b]MIT Operations Research Center and Laboratory for Financial Engineering; Lazard
[c]Gataca, Inc.
[d]MIT Sloan School of Management, Laboratory for Financial Engineering, and CSAIL; Santa Fe Institute
E-mail: alo-admin@mit.edu
[e]MIT Laboratory for Financial Engineering and Department of Electrical Engineering and Computer Science
[f]MIT Plasma Science and Fusion Center
*Corresponding author. E-mail: alo-admin@mit.edu

end of the century, the creation of a sustainable, unlimited energy source is of paramount importance. While solar and wind power technologies have improved in efficiency and cost, there are structural limitations preventing these renewable sources from supplanting the largely fossil fuel-dominated energy economy.

Fusion energy is a reaction in which energy is released by the *fusing* of atoms into new elements. This process, which powers stars and our sun, possesses the qualities of an ideal energy source: high-power density, high dispatchability (i.e., no intermittency), limitless fuel availability, a low environmental impact with high sustainability, and no risk of a runaway reaction, since it relies on its fuel remaining hot, not a chain reaction. And on December 5, 2022, the theoretical possibility of fusion became a reality when the Lawrence Livermore National Laboratory achieved net energy for the first time ever in a laboratory setting (i.e., the fusion reaction generated more energy than it took to initiate the reaction).[1]

The two main challenges in realizing this carbon-neutral future are the feasibility of fusion to produce net power at scale, and the uncertain timeline of near-term deployment of power plant demonstrations. While fusion research has historically been funded and executed by governments or government-sponsored projects, there has been a recent surge in private ventures aiming to embrace an approach different from the bureaucratic behemoth of ITER (which means "the way" in Latin), the world's largest fusion program, backed by the European Union, China, India, Japan, Russia, South Korea, and the United States. ITER has an expected budget north of $24B, promising to "produce 500 MW of fusion power," "... demonstrate the integrated operation of technologies for a fusion power plant," and "achieve a deuterium-tritium plasma in which the reaction is sustained through internal heating."[2] While ITER is generally regarded as a step in the right direction for fusion science, the US Department of Energy has cast doubt upon its projected cost, offering a significantly higher estimate of $65B (Kramer, 2018). ITER's timeline for its first plasma has faced delays, and there have been objections from the scientific community that ITER presents no direct pathway to economically viable fusion.[3] The cost point per megawatt of power would be one to two orders of magnitude away from that required to operate in a competitive energy landscape (NASEM, 2021).

On the private side of the fusion energy space, around 33 different companies are using a wide array of physics-driven techniques to explore fusion reactions. While the likes of Jeff Bezos, Bill Gates, and Peter Thiel have supported some of these ventures, the private fusion industry is still woefully underfunded compared to the progress of the research, as the cumulative capital raised is only around 20% of ITER's budget. For example, MIT's spinout, Commonwealth Fusion Systems, has displayed evidence that its tokamak reactor design, SPARC, can generate a 10-fold net energy gain (Creely *et al.*, 2020) based on access to more powerful superconductor electromagnets. Commonwealth expects to produce its first fusion plasma in 2025, and its compact design is projected to be a fraction of ITER's cost (Commonwealth has raised about $2B as of December 2022). The combination of high startup costs, an elevated risk of return for each individual venture, an atypical expected return horizon, and the sophistication of the underlying technology has generated lower capital deployment to the fusion space than what is warranted by its forecasted penetration into energy markets.

In this article, we synthesize the perspectives of present and potential future stakeholders in a fusion energy economy, including (but not limited to) academic research groups, government labs, startups, and energy-sector specific and general investors, to assess trends and views on private fusion financing, and then provide recommendations on how to leverage financial engineering methodologies to incentivize greater capitalization. During its compilation, we interviewed academics at leading research institutions; investors who have previously funded fusion projects; investors who have passed on fusion initiatives; members of energy advisor groups, investment banks, and sovereign wealth funds; founders of fusion startups; and members of the Department of Energy's Advanced Research Project Agency-Energy (DOE ARPA-E), the US government's research engine for energy technological exploration. The candor with which these stakeholders approached our interviews was evident (and appreciated), a testament to the vitality and zealousness of the fusion community.

We begin our discussion in Section 2 with an overview of the present energy landscape to motivate the need for a robust fusion energy industry. Pivoting to economics, in Section 3, we introduce current common attitudes toward financing fusion as determined by our interviews, along with an analysis of prior investment rounds in fusion. We present the Human Genome Project and the development of nuclear fission as case studies that may prefigure the potential evolution of the fusion sector. We then discuss auxiliary technologies that are pertinent to a successful fusion ecosystem and highlight the variety of spinoff applications that may result from core fusion research in Section 4.

Drawing on insights from our interviews, in Section 5, we propose a megafund securitization approach to financing fusion, in which many high-risk projects are amalgamated into a single financial entity, thus improving the risk–return balance of a portfolio of fusion projects to a point that funding might be sourced through a series of optimized debt securities in addition to an equity tranche. We further de-risk the megafund by leveraging first loss capital guarantees from philanthropic sources (ultra-high net worth [UHNW] individuals or private foundations) and governments to fund coupon payments to senior and mezzanine bondholders in its early years. This work represents the first application of the megafund concept outside of the biotech domain. This financial structure exploits the unique properties of the fusion sector, among them the ability of fusion companies to perform many types of divestitures (such as spinning off an independent company or executing equity carve-outs), the increasing industry demand for auxiliary technologies such as control systems, 3D printing, and rapid automated material testing, and the capacity for partnerships with academic institutions. This structure has the potential for immediate implementation, as long as the total assets under management (AUM) are consistent with the quantity of investable assets present in the fusion industry, but it will scale in relevance as the fusion industry proliferates in size and demonstrates commercial performance capabilities.

We then demonstrate performance simulations of the megafund vehicle by leveraging parameters of real fusion assets, estimated through collaborations with investors, academics, and large-scale philanthropic partners (Section 5.3). We conclude in Sections 6 and 7 with a discussion of the benefits and limitations of our megafund model, and outline future directions for our work. The Appendix provides a historical perspective on fusion research and practice, supporting mathematical formulations and derivations, and additional simulated results and sensitivity analysis.

2 The Current Energy Landscape

The 2019 World Energy Outlook Report, commissioned by the International Energy Agency (IEA), characterizes today's energy landscape as plagued by "deep disparities." Global carbon dioxide emissions totaled 36.8 billion tons in 2019, a 0.6% increase from the record-setting levels of the year before, despite the admonition in the recent National Climate Assessment report to reduce greenhouse gas emissions to prevent economic and infrastructural damage. Despite the pledges of numerous agencies to extend the reach of energy systems to developing and rural areas, 850 million people still lack reliable electricity, over 10% of the world population. Finally, the political ideal of an efficient and affordable renewable energy ecosystem is incompatible with the current reality of an economy dominated by oil and coal. In 2018, oil and coal accounted for 54.2 K and 43.7 K TWh (terawatt-hours), respectively, of global consumption, dwarfing the 3.3 K and 1.5 K TWh of consumption satisfied by wind and solar power.[4] This global reliance on oil is further complicated by geopolitical tensions and the elasticity of oil prices to world events. For example, global oil prices rose by 20% on September 16, 2019, following a missile attack on a Saudi Arabian oil facility.[5]

The disparities characterized by the IEA are expected to intensify in conspicuousness and effect, as an increasing population and strong economic growth, particularly in Asia, will require an expanding energy sector. The world's population is currently growing at an annual rate of 1.05%.[6] The US Energy Information Administration forecasts that global energy demand will increase by 50% by 2050. Under this projection of future energy production and consumption, three different scenarios have been assessed: a Current Policies Scenario, a Stated Policies Scenario, and a Sustainable Development Scenario.

2.1 IEA's future energy system scenarios

Under the Current Policies Scenario (CPS), the IEA contends that if energy dynamics persist on their present path, emissions will continue to rise while energy efficiency will stagnate, due to the absence of a strong financial incentive. As an alternative, the IEA's Stated Policies Scenario (SPS) enumerates the shifts in energy production and consumption if all countries abide by the existing and announced policies of the 2016 UN Paris Agreement and other programs. In effect, global energy demand would increase annually by 1% through 2040 (compared to 1.3% with the CPS), catalyzed by an expansion in the supply of low-carbon energy sources, chiefly solar photovoltaics (PV) and natural gas. However, under such a scenario, the consequences of growing energy demand would still far outweigh the benefits of the transformed cleaner-energy economy, and the sustainability objectives outlined in the Paris Agreement would not manifest themselves.

The final possible course of events is captured in the IEA's Sustainable Development Scenario (SDS), the predicted trajectory of energy sector dynamics if countries take the required action to fully realize the goal of the Paris Agreement to curb a global temperature increase to 1.5°C. Among the central policies embedded in the SDS are a gargantuan reliance on solar and wind power (with an increase in their net efficiency and a reduction in their costs), the modernization of energy systems with the aim of complete electrification by 2030, and incentivized shifts in demand toward efficient buildings, transport, and infrastructure. Like the SPS, the SDS is dependent on two main vehicles for its progress—international cooperation in government policy and economic incentives for renewable energy production—both of which face several challenges.

The principal difficulties in driving behavioral change through government policy can be found in the coordination of a uniform international policy, the mercurial (or corrupt) nature of national governments, and the power dynamics of federalism. Numerous studies have proposed game-theoretic models to induce cooperation among international agents with different resources and priorities (DeCanio and Fremstad, 2013; Wu *et al.*, 2014). In practice, it is nonetheless challenging to align the goals of nations. Given that carbon emissions are modeled as a public bad, the externalities of a defecting nation can be quite harmful to all. For example, the withdrawal of the United States from the Paris Agreement was expected to leave US emissions 3% higher in 2030 than previous estimates under the accord (Climate Action Tracker, 2019).[7] On the national scale, major hindrances to progress in climate policy include fluctuating party control in democratic societies, corruption in autocratic and oligarchic nations, the persistence of extractive states, and publicized misinformation (or a dearth of public education) on the present climate crisis (Lamb and Minx, 2020). Lastly, as with all matters of government, contention over federalism—that is, disagreements about the division of powers and responsibilities among different levels of government—presents a persistent challenge. In the US and Canada, competition between subnational governments over energy production and regulation has become routine, despite the lack of any national statute (Brown, 2012). As an illustration of this dynamic, in late 2007, consider that the governors of six states in the American Midwest (Illinois, Iowa, Kansas, Michigan, Minnesota, and Wisconsin) and the premier of Manitoba, Canada, signed the Midwestern Regional Greenhouse Gas Reduction Accord, an agreement to take governmental action to combat climate change. After the accord was signed, however, several of the Midwestern governors changed their stance on climate change and exited the coalition. The promised multi-sector cap and trade mechanism was never implemented.[8]

2.2 *Limitations in solar and wind power*

The challenges in developing a predominantly renewable energy sector derive from the technical and fiscal limitations of solar and wind power. It must be noted that recent technological and economic advances in solar cells are large and incontrovertible. In the past decade, the average power output of a 72-cell multi-silicon module climbed by 55 watts (around a 19% increase), while the unit price of a module declined sharply, from $2 per watt to $0.20 per watt, justifying a six-fold increase in installations during that period (Sun, 2019). Nevertheless, technical barriers remain, since the peak rate of efficiency for a PV module is only 23% (with its average falling between 15% and 18%). Solar is also inherently intermittent, and its time-averaged performance is further limited by its battery and inverter components (Timilsina *et al.*, 2014; IPCC, 2011). The laws of thermodynamics further restrain the performance of solar thermal applications, primarily the heat capacity of transfer fluids and heat loss during storage (Philibert, 2006). Economically, while solar panel prices have dropped significantly, and solar bonds have expanded the financial accessibility of solar power, its cost effectiveness is highly conditional on government subsidies, while its energy storage systems have yet to experience the same price reduction as its modules, and the required scaleup in energy storage is very large.

Wind power, meanwhile, has faced opposition on multiple fronts. Communities have opposed wind power due to the undesirable visual aesthetics of wind farms and their noise. Recent studies have also criticized wind power's poor average power density and the quantity of energy produced relative to the physical area of the turbine, and they question the feasibility of large-scale deployment of wind power. For instance, the largest wind plant by capacity, China's Gansu

Wind Farm, comprises 30 square kilometers of land, half the size of Manhattan, but at peak only powers the equivalent of 60% of Manhattan's retail demand.[9] Besides its well-documented problem of intermittency, since wind power is subject to temporary stoppages due to exogenous conditions, Miller and Keith (2018) note the existence of a "wind shadow" effect, in which the close proximity of spinning blades creates drag in the wind. Adding more turbines to a farm will thus lower the power output of each individual unit. Wind power also faces an additional institutional barrier, since novel wind power equipment has poor compatibility with existing energy system infrastructure compared to fossil fuels (Herbert *et al.*, 2014).

Overall, scientists have emphasized the necessity for a dispatched central carbon-free energy source to supplement wind and solar to achieve complete decarbonization feasibly from a cost perspective. A 2018 study out of MIT evaluated two paths for decarbonization: one that exclusively permitted solar and wind generation with battery storage, and another that additionally considered continual carbon-free sources, such as fusion, bioenergy, fission, and natural gas embedded with a carbon capture mechanism (Sepulveda *et al.*, 2018). It is concluded that the cost savings of the mixed energy sector were substantial, with a peak reduction of 62% in one scenario.

But this need for a robust renewable energy sector, and furthermore a heterogeneous renewable energy sector, has not manifested in government research spending. Historically, government and institutional uncertainty about whether a future energy system can depend solely on renewable sources has induced a disproportionate allocation of research and development resources to fossil fuels and energy efficiency technologies. Between 1978 and 2018, US federal spending on renewable energy R&D accounted for only 18% of the Department of Energy's R&D budget, marginally higher than the 16% earmarked for energy efficiency (Clark, 2018).

To chart a course toward an eventual carbon-free energy system as meticulously as possible, it will be imperative to augment current methods and research on sustainable energy with other forms of energy innovation. Fusion energy may fit the niche that other renewables and fission have struggled to fill.

3 Fusion Financing and Private Nuclear Fusion Companies

Historically, capital- and time-intensive projects have had difficulty finding sufficient continuous financial support. This is especially true for initiatives that typically yield their profits (if any) over a 20-year timespan. This investment profile is generally unattractive for private investors and venture capital firms, which typically seek returns within 3 to 5 years. Consequently, governments have historically funded fusion research projects. Fusion funding in the US peaked during the energy crisis of the 1970s, sharply decreased in the 1980s, and has since remained relatively constant (Figure 1). The total cumulative

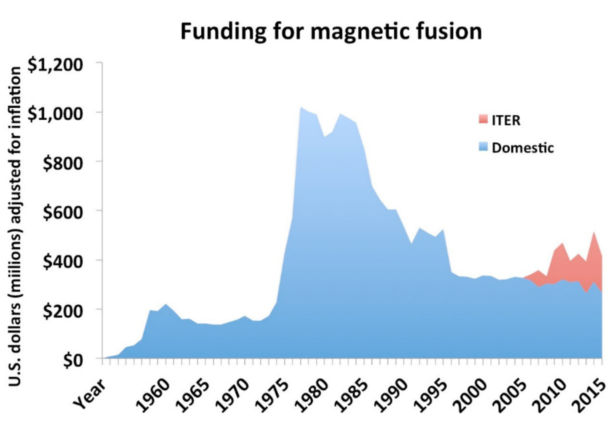

Figure 1 Time series of funding data for magnetic fusion. Data provided by the US Department of Energy.

US spending on magnetic fusion confinement is nearly $18B, adjusted for inflation.

The impetus behind creating a global consortium to fund ITER was for nations to share the risk of fusion development. However, the added complexity of a project funded and constructed by multiple countries, their subsequent political disagreements, and a budget that has ballooned by a factor of 10, has greatly diminished the public funding opportunities for new projects in the US. The US alone has contributed $1B to ITER through 2018, with plans to give another $500M through 2025.[10] With its pledge to contribute approximately 9% of the construction cost (mostly in kind), the US contribution to ITER in its construction phase is estimated to be nearly $5B. It is thus difficult for most governments to justify the construction of localized reactors in ITER-participating countries.

The two exceptions are China and the UK. In China, the Ministry of Science and Technology announced in 2011 plans to construct CFETR, a new tokamak similar to ITER, which will initiate operations by 2030, to complement its existing EAST initiative. China is one of the few governments capable of securing another major investment in fusion, as China will account for 30% of the increase of the total world's energy demand between 2014 and 2040 (IEA, 2019). The UK likewise announced its STEP program, focused on being the first nation to deliver fusion energy to the grid by 2040. We should also note that in 2020, there was increased discussion among scientists in the US, in conjunction with the Department of Energy's Fusion Energy Sciences program, to operate a pilot plant on American soil by the 2040s.[11] In February 2021, the National Academies of Sciences issued a consensus study report to encourage public–private partnerships to develop a robust fusion ecosystem in the US. The report's objectives are aimed at developing practical energy production with structured costs and marketplace penetration. Its timeline calls for the greater fusion community to achieve "viable design by 2028 and initial pilot plant operation in 2035–2040." The report's stated goals possess enough generality to apply to any core fusion concept, but are best suited for a deuterium–tritium fueled tokamak, identified in the report as the "leading fusion concept."

For countries other than China, the exclusively public funding of nuclear fusion projects is not feasible. However, there is a growing alternative in private capital. Private fusion financing is still in its infancy, as the total capital invested in this space is only $4.8B through 2022 (Fusion Industry Association, 2022). A majority of the sector's capitalization can be accounted for by its four largest private ventures: TAE Technologies, Commonwealth Fusion Systems, General Fusion, and Tokamak Energy (see Section 3.3).

3.1 Comparable industries

To forecast the evolution of the fusion energy sector, we briefly examine two comparable industries that possessed analogous initial technical hurdles and grew because of both public and private investment: nuclear fission and the Human Genome Project.

3.1.1 Nuclear fission

The breakthroughs in research that led to the creation of nuclear fission power were financed primarily by the US federal government. This research originated as a direct result of World War II, with the practical objective of creating an atomic bomb. The federal government funded the Manhattan Project through the Army Corps of Engineers, spending $2.2B. After the war, Congress passed the Atomic Energy Act, which transferred nuclear research from the military to a civilian government agency, the Atomic Energy Commission (AEC). In 1951, the AEC constructed the first experimental nuclear reactor to generate electricity in Idaho. The US Navy played an instrumental role in the research and design of nuclear fission reactors as part of its development of nuclear-powered submarines. Alvin Weinberg of the Oak Ridge National Laboratory helped develop the light-water reactor, which is the most common type of reactor in use today (Weinberg, 1994).

While government laboratories accomplished most of the research breakthroughs necessary for nuclear fission, the construction of nuclear plants was financed by a combination of public and private financing. As will likely be the case with nuclear fusion, erecting a fission-based power plant requires enormous upfront costs and a multi-year construction timeline. The first full-scale electrical atomic plant, Shippingport, was built in 3 years for $72.5M (approximately $650M in 2020 dollars), opening in 1958. In the 1960s, the majority of nuclear plants were funded purely by governments, but a few plants were fully financed by the private sector. Nuclear technology had matured by the 1970s to compete for market share through electricity generation. During this period, governments continued to support nuclear plant development both directly, through financing, and indirectly, by regulating energy markets (thus providing utilities price certainty) and offering guarantees. Purchase power agreements, contracts in which buyers agree to purchase electricity at a predetermined price and quantity over a certain term, were often insured by governments.

Private financing of nuclear power nevertheless faced obstacles, as modeling a return on investment still proved difficult in light of the long-term nature of the investment and the fluctuating prices of electricity if the plant should be located in an unregulated electric power system. Additionally, many private investors became dissuaded by frequent construction delays, cost underestimates, and the occasional revocation of permission to build already partially erected plants. Private sector investment in US nuclear power plants essentially froze beginning in the late 1970s, due to the combination of the high cost of nuclear plant construction, the Three Mile Island accident (a partial meltdown of a reactor in Pennsylvania), and energy deregulation. There has been only one new reactor built in the US since 1996, when a reactor at the Tennessee Valley Authority's Watts Bar nuclear plant was completed in 2015, although two reactors at the Vogtle power plant in Georgia are still under construction.

To counter these obstacles in nuclear construction, other countries have implemented innovative private financing models, Finland being a notable example. Finland employs the cooperative Mankala model, in which investments in nuclear plants are pooled among a group of companies. Shareholders have the right to purchase electricity from the plant proportionally to their stake, which can either be used or sold.[12] France has used a similar model, in which a large number of French companies formed a consortium called Exeltium, which then entered into a contractual agreement with the operator of France's nuclear power plants, Électricité de France, to provide upfront financing in exchange for a long-term contract to supply electricity at favorable prices.[13] A third model comes from the UK, whose government in 2014 introduced contracts for difference, in which a third-party investor agrees to bear responsibility for the difference between the cost of the plant (and a predetermined profit margin) and the market price for electricity. Under such an arrangement, if the price of electricity falls below predicted amounts, the third-party investor would cover the cost, and if the price of electricity exceeds predicted amounts, the third-party investor would reap the subsequent profit.

3.1.2 Human Genome Project

Government-supported funding also played a significant role in another important scientific breakthrough of the twentieth century, the decoding of the human genome (Gisler *et al.*, 2010). In 1990, the US Department of Energy and the National Institutes of Health agreed to a memorandum to provide over $3B in funding for the Human Genome Project (HGP), which was expected to take 15 years to complete. The advances in medicine and science that were catalyzed by the HGP have generated nearly $800B of additional economic activity in the US, making the government's investment one of the most profitable in US history.[14] Celera, a private company, also invested $300M to decode the human genome, on a parallel track to the HGP. However, when President Clinton announced that the genome could not be patented, Celera was not able to fully monetize its investment.

While Celera still played an important role in decoding the genome, this cautionary tale highlights the tension between private investment in major scientific initiatives and the free sharing of information to stimulate further scientific advances. However, unlike nuclear fission, the government's funding of the HGP has spurred many tens of billions of dollars of private research in the areas of medicine, pharmaceuticals, agriculture (including livestock), and other industrial processes.[15]

3.2 Overview of the private fusion energy space

Recent technological developments have caused fusion research and reactor construction to begin their transition from government laboratories to the private sector. There are 33 active fusion companies and 31 in the Fusion Industry Association, a nonprofit organization composed of companies working to commercialize fusion power.[16] Demonstrating the close link between the private sector and academia in this space, over half of listed fusion companies are associated with or spun out of a university setting. Advisory organizations in this space have also formed, e.g., the not-for-profit Stellar Energy Foundation, under the leadership of Jesse Treu, which has been actively consulting with fusion startups and assisting with fundraising.

As depicted in Figure 2a, private fusion companies began to proliferate in the early 2010s, and this trend accelerated after 2017. This corresponds with an increase in private-sector investment in fusion, which doubled in 2016 and was generally higher throughout the 2010s compared to the previous decade, as shown in Figure 2b. The current total financing in this space is a little over $4.8B, as UHNW individuals such as Jeff Bezos, Bill Gates, and Peter Thiel have thrown their support behind such ventures as well as institutional investors and energy companies. It should be noted that a very large percentage of investment in this sector went into the four largest startups: TAE Technologies, Commonwealth Fusion Systems, General Fusion, and Tokamak Energy.

In analyzing the startup space, it is important to consider the heterogeneity in startup business models in addition to their technological variation. Figure 3 catalogs several potential directions for fusion companies, with the main strategies being: (1) developing intellectual property (IP) and licensing it to government-led fusion programs, other private fusion startups, or energy companies; (2) designing, building, and selling reactors to another party; or (3) constructing power plants to actively compete in electricity markets. These business methods require different amounts of resources and capital, and therefore can appeal to different types of investors

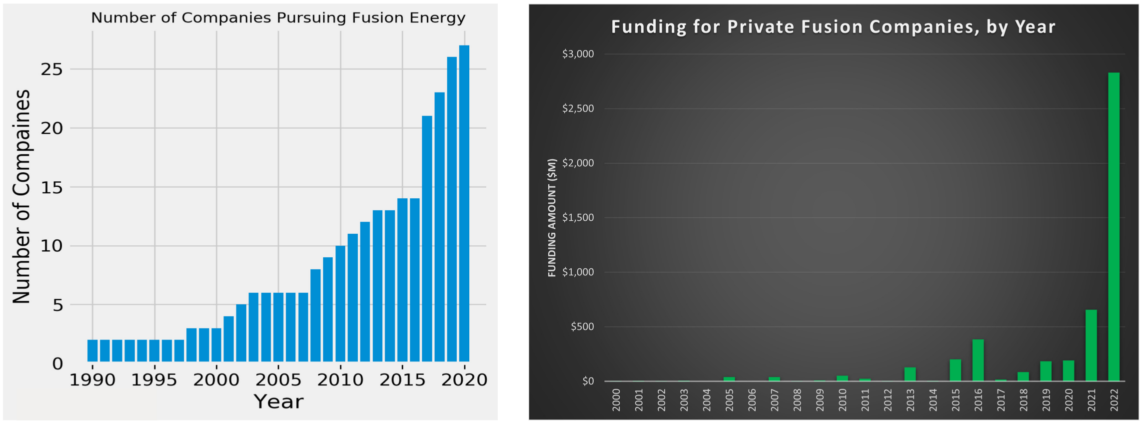

(a) (b)

Figure 2 (a) Growth in the number of private-sector fusion companies; and (b) private-sector financing of fusion.

Sources: (a) https://www.fusionenergybase.com/article/the-number-of-fusion-energy-startups-is-growing-fast-heres-why; (b) https://www.fusionenergybase.com/article/funding-to-fusion-energy-companies-since-2000.

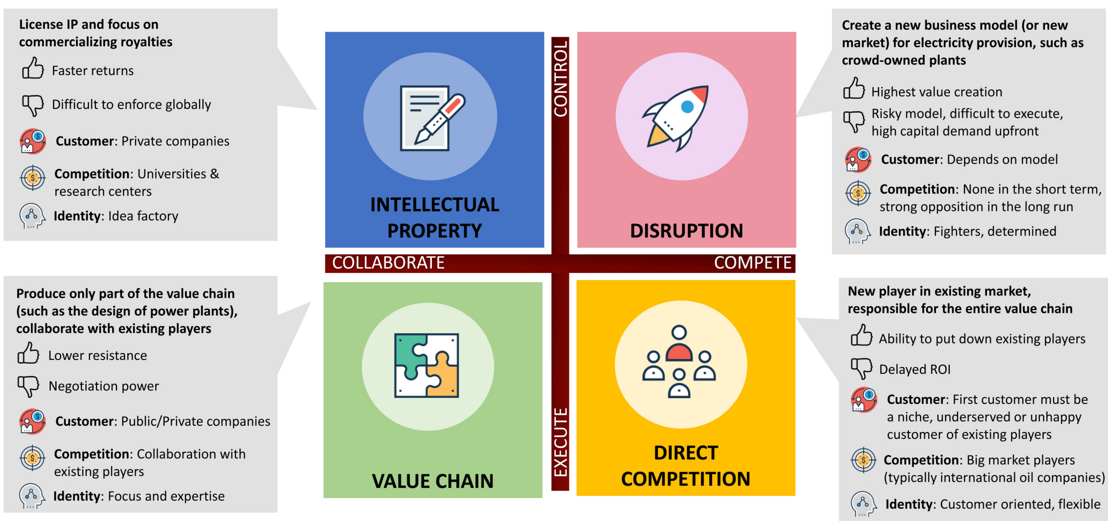

Figure 3 Business models for fusion energy companies.

with varying risk appetites. Private ventures using all three business approaches have seen some success in the nuclear fission space.

3.3 *Key players*

With these business models in mind, we now provide brief overviews of the largest private fusion companies currently operating.

3.3.1 *TAE Technologies*

TAE Technologies emerged from the University of California, Irvine in 1998. It is based in Foothill Ranch, California, and has raised about $1.2B since inception from notable institutional firms such as New Enterprise Associates and Venrock, sovereign wealth funds, and high net wealth individuals, including Charles Schwab and John Mack. Using a field-reversed configuration, TAE intends to use hydrogen–boron fuel and insert high-energy hydrogen particles for its fusion reaction, which would produce minimal neutrons, an important design consideration. (While the primary fusion reaction under investigation by TAE is aneutronic, a small percentage does form neutrons due to secondary reactions.) TAE has schematics for three main reactor designs of varying size: Norman, a test device currently in use; Copernicus, which is in early development, and is intended to demonstrate breakeven energy; and Da Vinci, a power plant-scale reactor that will be constructed in the late 2020s.[17]

3.3.2 Commonwealth Fusion Systems

Commonwealth Fusion Systems spun out of MIT's Plasma Science and Fusion Center (PSFC) in 2017. Based in Cambridge, Massachusetts, it has raised over $2B in capital since inception from a cohort of venture capital firms, including (but not limited to) Breakthrough Energy Ventures, The Engine, and Khosla Ventures, as well as the Italian oil and gas company ENI. Commonwealth's compact reactor design, SPARC, makes use of novel high-temperature superconducting magnets. Commonwealth has recently built a coil capable of producing a magnetic field of 20 tesla, and aims to demonstrate breakeven energy by 2025.[18]

3.3.3 General Fusion

General Fusion was founded in 2016. Located in Burnaby, British Columbia, it has raised about $300M in funding from private sources, including Jeff Bezos' venture vehicle and the Canadian oil and gas company Cenovus Energy, as well as from public sources, in this case, the Canadian government. Its core technology uses magnetized targeted fusion, and it generates breakthrough energy with a compression ratio of 7:1, compared to inertial confinement's standard ratio of 40:1. General Fusion is currently building a reactor to be completed in 2027, and with the data gathered from the device, it will begin developing a power plant-scale reactor shortly thereafter.[19] In June 2021, General Fusion announced that it would build its Fusion Demonstration Plant at the UK Atomic Energy Authority's Culham Campus.[20]

3.3.4 Tokamak Energy

Tokamak Energy, formed in 2009 in Arbingdon, Oxfordshire, has raised about $250M of financing since inception as of October 2022, which includes investments from private individuals and the UK company Legal and General, as well as government grants and research and development credits. Tokamak Energy's current device, ST-40, is a spherical tokamak that uses standard magnets. By incorporating high temperature superconducting magnets into its next design, to be built by 2026, Tokamak Energy hopes to exhibit breakeven energy.[21]

3.3.5 First Light Fusion

First Light Fusion was spun off from the University of Oxford in 2011. Based in Yamton, Oxfordshire, it has raised approximately $107M to date with lead investor IP Group. First Light is using a version of inertial confinement that employs a railgun-type approach to stimulate a fusion reaction. Its timeline projects that its first power plant will be built in the 2030s.[22]

3.3.6 Helion Energy

Helion Energy grew out of research conducted at the University of Washington in 2013. Based in Redmond, Washington, it has obtained $578M in financing from investors since inception including the Thiel-founded Mithril Capital Management and Y Combinator. Helion's device uses magneto-inertial confinement with a field-reversed configuration, which is eventually intended to use a unique fuel combination of deuterium and helium-3. Helion has developed a prototype and is working toward showing breakeven energy.[23]

3.3.7 *Zap Energy*

Zap Energy spun out from the University of Washington in 2017. Located in Seattle, Washington, it has raised $202M to date with investors including Chevron's venture arm and Lowercarbon Capital. Zap uses a modified form of the Z-pinch called the sheared-flow stabilized Z-pinch, in which the velocity and radii of plasma flow are variable. After promising results in increasing the flow of electrical current, Zap aims to achieve breakeven energy in as little as three years.[24]

3.4 *Fusion investors*

To further assess the state of private fusion, we formed a Fusion Energy Investor Database—compiled through public disclosures from companies and investors, data from PitchBook, data from Crunchbase, and information from investor interviews—to track all private investment rounds and analyze the participating stakeholders. Figure 4 displays some key characteristics of

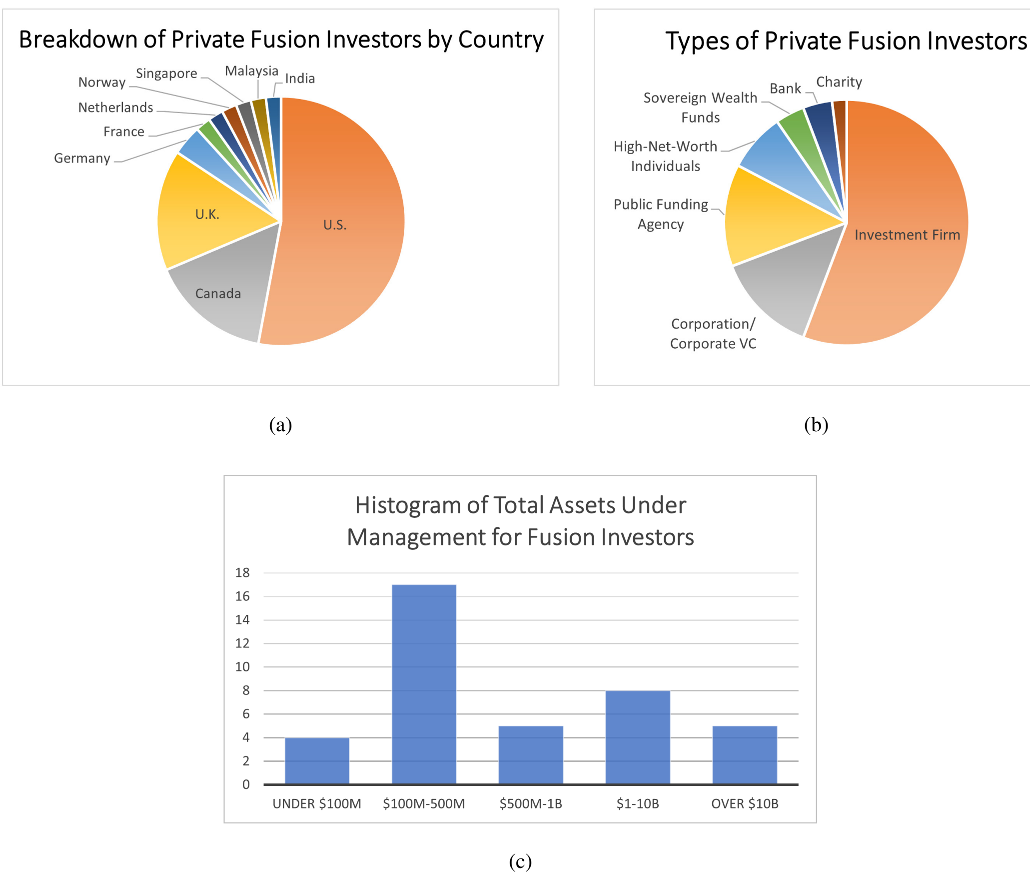

Figure 4 (a) Proportional headquarters of current fusion investors; (b) the division of fusion investors by type; and (c) a histogram across all past fusion investors of total assets under management as of 2020. For corporate investment, we take this value to be the size of the corporate venture fund or sum of all investments made, not the total valuation of the corporation itself. For individual investors, we use net worth.

the fusion investor community today. The majority of investors in this space are based in the US, followed by the UK and Canada, which is not surprising, given that all the major private fusion ventures are in these countries. Additionally, investment firms (including venture capital, private equity, and family-run offices) tend to dominate the initial funding of fusion. Corporations, specifically those in the energy sector, have been active as well, accounting for 70% of the cohort. For example, the Canadian oil and natural gas company Cenovus has invested in General Fusion, the Italian oil and gas company ENI and the Norwegian energy company Equinor participated in Commonwealth Fusion Systems' raises, and most recently, the US energy giant Chevron funded Zap Energy. This may indicate a trend that energy companies are beginning to consider fusion as a future sector for expansion. To the best of our knowledge, sovereign wealth funds from two countries (Malaysia and Singapore) have thus far contributed to fusion ventures. Regarding the size of investment firms, the majority has AUM in the \$100M to \$500M range, which is standard for early-stage venture capital. We anticipate that this mode will shift upward in the coming years, as fusion companies seek larger round sizes, and investment in this space becomes more mainstream. Investment companies with higher AUM will be better suited to continuously support the growth of businesses, both financially and functionally, as they migrate from ideation and development into production and operation.

3.5 Financing attitudes

For all investors, the canonical objective is to maximize their returns at minimum risk. Before we can assess how an opportunity to invest in fusion fares under this function, we need to discuss three fundamental aspects of fusion projects that must be taken into consideration.

First, fusion companies are estimated to require a substantially larger amount of upfront investment than traditional early-stage tech companies, even for early-stage R&D (climate tech in general tends to induce larger venture capital deal sizes, as a little under two-thirds of these deals exceed \$100M[25]). This implies that only large venture capital firms and wealthy private funds will be able to lead deals, thus narrowing the number of options to high net worth investors, large sovereign funds like the Oil Fund of Norway, multinational holding companies like SoftBank, and large venture capital funds like Sequoia Capital or Energy Venture Capital.

Second, the main goal of many fusion scientists is to develop an affordable source of clean energy for the world. They are largely motivated by altruism, rather than purely by economic aspirations. This perspective can clash with investor incentives, which poses a risk in sustaining economic support for the project and the ability to make choices in a desirable direction.

Third, unlike most new ventures, non-government fusion projects are typically incubated in larger research organizations or universities. These structures provide the necessary economic and human resources to develop the relevant IP to bring the project to fruition. For the venture to scale and provide measurable value to its shareholders, it is advisable to create an independent legal entity or spinoff that holds exclusivity in these resources. This step introduces an additional axis of complexity, as the parent organization and the spinoff need to engage in potentially complicated negotiations that could affect an investor's risk assessment of the project.

With this in mind, we conducted interviews with a range of investors—some who have previously invested in fusion startups, some who are small and medium enterprise investors in energy, some who work at sovereign wealth funds, some who work at traditional venture capital funds, and some who had never heard of fusion energy prior to our conversation—and present the following qualitative analysis. Those people surveyed were promised anonymity to maintain the integrity of their disclosure. We do not tabulate these qualitative insights into a quantitative data form due to the statistically insignificant sample size. We further wish to qualify that, while all viewpoints presented were expressed by multiple individuals, they may not apply to all potential investors in the energy or fusion sector.

(1) **Traditional venture capitalists are not knowledgeable about the technical specifics of fusion.** While this is to be expected, given the niche state of fusion, several investors admitted to executing an investment without sufficiently comprehending the nature of the technology. That said, there are sector-specific venture firms that do possess a deep knowledge of fusion research, such as Breakthrough Energy Ventures and the fusion-specific fund Strong Atomics.

(2) **Investors struggle to predict the probability of success and future financials of fusion companies.** Many investors claimed to have tried typical valuation models in the fusion space to no avail. This results in fusion companies being evaluated relative to later-stage energy startups, or investors excluding fusion opportunities from their investment screening process.

(3) **Many institutional funds are not interested in opportunities with fusion's time frame (10+ years).** Some investors asserted that despite potential interest in a fusion investment, their fund structure is not conducive to investing in a venture that will not show any potential cash flows for a minimum of 10 years.

(4) **Most fusion investment opportunities to date are "too early stage" for investors**. Several investors expressed a willingness to invest in fusion startups that were in the process of prototyping or developing a power plant, but this only applies to a handful of current ventures. The majority of investment opportunities presented is for companies to engage in early R&D.

(5) **Investors value technological and IP differentiation.** As is the case with most sectors, investors highly value differentiation from existing competitors. The causal effect of this preference encourages variety in technological approaches for startups, which is positive for the sector. However, it could also cause infeasible or less tenable concepts to be overfunded relative to those grounded in historical research.

(6) **Investors seem to have a preference toward magnetic confinement over inertial confinement.** While this may be a product of more startups using magnetic confinement, some investors have cautioned that inertial confinement may be more geared toward military applications than energy production, since federal funding for inertial fusion arises primarily from the National Nuclear Security Administration.

(7) **Many investments to date have had altruistic philanthropic underpinnings.** Many angel investors, and high net worth individuals in particular, have signified that their investment into fusion startups was not motivated by profit. To this end, some investors have made personal investments in fusion rather than going through their institutions.

(8) **Fusion could be a long-term hedged investment for investors that are long oil.** Sovereign wealth funds, large venture capital, and private equity funds can use investments in fusion as a strategic tactic to

diversify against the risk of oil depletion and increased government regulation. However, this stratagem can only function if investors are able to go long on the fusion industry as a whole, not merely for an individual startup or project.

(9) **Once breakeven energy is achieved, further evaluation and capitalization of fusion companies will balloon.** Some investors with familiarity in other sectors compared fusion financing with the early funding trends of Internet and biotech companies. The most compelling illustration was how the initial public offering (IPO) of Genentech in 1980 (in which the price of the stock doubled) catalyzed a slew of biotech IPOs and generated much greater investor interest in the biotech field.[26]

(10) **Investors fear underbudgeting and overly ambitious deadlines ("ITER Effect").** Investors familiar with the fusion space have pointed to the logistical shortcomings of ITER to justify their hesitation in investing private capital in startups. The same trend has also prompted some investors to play a more active role in their investment, to ensure that companies are adhering to their proposed budget.

(11) **Investors are skeptical about opportunities in the space that seem too good to be true ("Theranos Effect").** The technical complexity, large overhead costs, and often secretive nature of fusion seem to have generated an association with the fraudulent Theranos company. Investors have now made a point to analyze a firm's core technology rather than overvalue assets because of prominent board members or advisors to the company, or unsubstantiated lofty promises. In fact, startup executives and investors alike now fear that the failing of any private fusion company could disincentivize further investment in the space.

4 Auxiliary Companies and Fusion Spinoffs

Thus far, we have only examined core fusion technologies used to develop energy-producing reactors. Before we discuss possible financial structures that satisfy the aforementioned investor attitudes, we highlight two additional unique company categories that could be incorporated into a fusion fund: service and technical companies that would be needed in a fusion economy to supplement fusion reactors (auxiliary companies), and applications of fundamental fusion technology that can be monetized in separate industries (spinoffs). Auxiliary companies are essential to the advancement of a modularized fusion industry, and they offer a different risk–reward profile than core fusion assets. Spinoffs allow for fusion startups to generate much nearer-term cash flows, and they may lead to significant technological advancement in certain fields. We also review the current state of fusion powered rocket engines.

4.1 Auxiliary technology and research

As with any burgeoning industry, there are a number of related auxiliary technologies that will be necessary to sustain a future fusion ecosystem. In our interviews with academic researchers and executives at fusion startup companies, many of our subjects expressed the advantages of introducing a high degree of modularity into the industry, specifically to simplify the development pipeline and supply chain for companies building reactors. We have identified several existing and new areas that may warrant increased innovation and investment, including (but not limited to) tritium fuel supply, handling and storage, control systems, specialized materials manufacturing (two examples are vacuum vessels and first wall components), reactor maintenance, generator and heat exchangers, lithium cooling blankets, superconductor suppliers, and 3D printing. Although some of these auxiliary components are presently being researched at academic institutions, we foresee the potential for near-term commercialization in many of these areas. Below, we describe several auxiliary technologies that are in various stages of commercial development.

4.1.1 Fuel

Most fusion reactions currently under investigation in the fusion space require deuterium and tritium to generate energy. While deuterium is a common isotope of hydrogen, abundant in water, tritium is generated in CANDU-type fission reactors through the interaction of fission neutrons with a heavy water moderator, and the resulting tritium is extracted from the heavy water moderator by a tritium removal facility (Phillips and Easterly, 1980). Only three tritium removal facilities exist in the world (in Canada, South Korea, and Romania). Those three facilities do not produce enough tritium to support the requisite number of fusion plants. To provide enough tritium to fuel many fusion reactors, "breeding blanket" technology will need to be developed and employed, in which neutrons from a fusion reaction react with lithium in the blanket to "breed" tritium fuel to continue to power the reactor. The creation of a lithium-breeding blanket is still in the research stage.

4.1.2 Control systems

Fusion requires a plasma control system to regulate the plasma through algorithms based on measurements from thousands of diagnostic signals. General Atomics designed the plasma control system used by ITER, and their researchers at the DIII-D National Fusion Facility have used machine learning to estimate the change in instability of the plasma. Princeton University has also developed advanced machine learning software, known as the Fusion Recurrent Neural Network, for use as a predictive tool by ITER.[27] Researchers at several universities, such as the Lehigh University Plasma Control Group, are focused on optimal control systems for stabilizing plasma.[28] The general digital control systems market has been forecasted to be worth $25.82B in 2026.

4.1.3 Materials science

A key challenge in fusion research is developing plasma-facing materials that will not degrade from being in proximity to the fusion plasma. Deuterium–tritium fusion reactions produce high-energy neutrons, which collide with atoms in core materials and induce a cascade. This reaction can damage the plasma-facing wall, vacuum vessel, certain electronic components, and even magnets. As a result, using damage-resistant materials in the reactor is imperative. Additionally, tokamaks and stellarators require materials that can sustain a high magnetic field. For this reason, Commonwealth Fusion Systems is using a high-temperature superconductor cable made from rare earth barium copper oxide (REBCO; see Section 4.3.2 for further details). REBCO was initially produced by two companies, Superpower and Bruker Corporation (Romanov *et al.*, 2020), and is now commercially available through a handful of companies.

4.1.4 3D printing

3D printing technology has numerous potential uses in fusion. For example, scientists at China's Shenzhen University have produced tritium breeding blanket components using 3D printing that are more reliable and efficient than those produced through traditional processes (Cui *et al.*, 2020). At the Institute of Nuclear Energy Safety Technology in China, scientists are using a form of 3D printing called selective laser melting to fabricate additional components of a fusion reactor.[29] The total market capitalization for 3D printing was $13.84B in 2021, and it is expected to grow quickly to $37.24B by 2024.

4.1.5 Lithium cooling blanket

In addition to "breeding" additional tritium fuel, a lithium blanket in the fusion reactor will also serve as a cooling mechanism by absorbing energy from the neutrons produced within the plasma of the fusion reaction. Numerous research laboratories and private companies are working on designs to increase the effectiveness of a lithium blanket. Assystem's Sunderland engineering team was awarded the contract for the lithium blanket design for the United Kingdom's Spherical Tokamak Energy Production, which is a government-financed project to build a commercial fusion energy power station by 2040.[30]

4.2 Fusion energy rocket engines

While space travel does not qualify as an auxiliary technology to the fusion energy space, its future may be dependent on powering spacecraft by fusion-powered rockets. The use of fusion would provide more efficient propulsion than chemical combustion, and it would allow rockets to travel at much higher speeds. NASA's Fusion Driven rocket uses a solid lithium propellant, requiring no significant tankage mass, and enabling a rocket to reach Mars in 90 days, compared to 8 months using a conventional rocket.[31] One significant challenge in developing a fusion-propelled rocket is that the fusion reactor must be compact. Additional concerns include the functionality of thrusters under extreme atmospheric conditions, and the energy inefficiency of current ion thrusters. NASA has funded research into potential designs conducted by the University of Washington and the private company MSNW LLC. In addition, the Princeton Plasma Physics Laboratory is working on a project in conjunction with the spinoff Princeton Satellite Systems, also under a NASA grant.[32] The nonprofit Limitless Space Institute, established in 2019 by astronaut Brian "BK" Kelly and Sonny White, the former head of advanced propulsion systems at NASA, has advocated for a fusion energy consortium to develop fusion propulsion technologies to enable interstellar travel. In the private sector, Apollo Fusion has been acquired by Astra to support its efforts to develop a fusion rocket.

4.3 Spinoffs for fusion companies

Another crucial consideration when evaluating a fusion investment is other uses for its core research and IP. The underlying technologies of fusion, including plasma physics, materials science, advanced computational methods, and nuclear engineering, have potential innovations in numerous other industries, including medical technology, pharmaceuticals, electrical delivery, robotics, and superconducting technology. Although the return on investment for fusion as an energy source may take a decade or longer to achieve, the laboratories and private companies involved in fusion research can, and likely will, use their technologies in subsidiary commercial endeavors that will produce shorter-term returns. Because many technologies involved in fusion have the potential to innovate in industries with a high economic impact, the value of the spinoff technology must be factored into the return on any investment in a fusion startup. Some areas in which fusion research already has been applied to other industries are described below.

4.3.1 High-temperature superconducting magnets

A significant area of fusion research is the development and application of high-temperature superconducting magnets to generate a magnetic field stable and large enough to confine the plasma. Compared to previously used low-temperature superconducting magnets, high-temperature superconducting (HTS) magnets can be constructed much smaller and can generate substantially larger magnetic fields. Commonwealth Fusion Systems has developed a 20-tesla HTS magnet for their pilot tokamak.[33] The applications for HTS magnets in other fields are numerous, as they offer a significant cost reduction and performance improvement compared to traditional magnets. Certain spinoff markets include (but are not limited to) nuclear magnetic resonance, MRI technology, generators in wind turbines, superconducting energy storage, magnetohydrodynamic generators, and material separation in mining. Some specific areas of applications are further detailed in this section.

4.3.2 High-temperature superconducting technology

Fusion startups are also developing other HTS technologies that can seamlessly integrate into novel superconducting magnets. In October 2020, an MIT-led group of scientists reported their design of an HTS cable named VIPER, which was described as a "vacuum pressure impregnated, insulated, partially transposed, extruded, and rolled-form cable" (Hartwig *et al.*, 2020). The VIPER cable, which contains yttrium barium copper oxide, is used by Commonwealth Fusion Systems, and has the potential to transform numerous industries. For example, the power industry uses copper lines to transmit electricity in the power grid. By replacing the lines with HTS cables, which provide less resistance to the flow of electricity than conventional copper lines, with five times as much electric current, electrical transmission losses can be reduced by half (from about 8% to 4%) (Masuda *et al.*, 2007). In addition to these cables, manufactured by the Bruker Corporation, the Oak Ridge National Laboratory has developed its own superconducting cable in partnership with the Southwire Company. The projected market size for HTS technology was estimated to be $488.6M in 2020, and is projected to grow to $546.7M by 2026.

4.3.3 Cyclotron advances

Cyclotrons accelerate particles to high energies using confining magnetic fields. They have numerous uses in scientific research, with an additional medical application to treat certain forms of cancer with proton beam radiotherapy (PBRT). Proton beams can be customized to the size and shape of cancer tumors, thus causing less damage to healthy tissue than traditional radiotherapy (Gales, 1999). However, a PBRT cyclotron is extremely large, around 100 tons, and costs approximately $100M to construct.

The superconductivity research performed in connection to fusion development may provide the technology to develop a superconducting cyclotron that is smaller and more affordable than the PBRT cyclotrons presently in use. For example, MIT's PSFC is developing a compact synchrocyclotron that is about 40 times smaller and lighter than a traditional PBRT, about 25 tons in mass.[34] The global market for radiotherapy for cancer is approximately $7B, and it is projected to reach $10.15B by 2025. While proton therapy constitutes only a small percentage of that market, the development of less costly superconducting cyclotrons may result in proton therapy capturing some of traditional radiology's market share.

4.3.4 Electromagnetic aircraft launch system

General Atomics, a San Diego firm with heavy involvement in ITER, has applied its fusion research in electromagnetics to develop a replacement for the high-pressure steam catapults used to launch planes from aircraft carriers.[35] Through its Electromagnetic Aircraft Launch System (EMALS), an electromagnetic catapult will replace the steam catapults used on earlier generations of aircraft carriers. The technology will reduce costs and allow fewer staff to be involved in the launch. The US currently operates 11 aircraft carriers, all of which may eventually use EMALS technology.

4.3.5 Thermal measurement in extreme conditions

Fusion research into plasma monitoring under conditions of extreme heat is also applicable to the production of metals, glass, and other materials that are produced under extreme temperatures. For example, the MilliWave Thermal Analyzer, developed by MIT's PSFC in conjunction with Pacific Northwest National Laboratory and Savannah River National Laboratory, can monitor the properties of materials inside a glass melting furnace or a process reactor.[36] The thermal analyzer uses millimeter-wave electromagnetic radiation to constantly monitor the temperature and chemical properties of materials, thus eliminating the need for workers to take manual measurements. The global market for thermal analyzers is expected to reach $670M by 2028.

4.3.6 Robotics for remote maintenance in harsh environments

Fusion researchers are developing robots that can withstand the extreme conditions inside a nuclear fusion reactor. The UK firm Assystem Engineering has built robotic equipment that can safely remove irradiated components from tokamaks, which it will test with ITER. The UK Atomic Energy Authority has a research facility for remote applications in challenging environments (RACE). The focus of the RACE facility is the design of robotics and artificial intelligence systems that will install and maintain equipment inside a nuclear reactor.[37] This technology has obvious applicability to numerous other areas, including natural disaster response, bomb disarmament, search and rescue situations, military operations, and firefighting. The artificial intelligence developed to navigate a robot inside of a fusion reactor is also applicable to the autonomous driving sector. According to a 2018 McKinsey Report, the market for industrial robotics has been growing at a double-digit pace, and is valued at $48B in the US alone.

4.3.7 Medical scanning devices

HTS magnets developed for nuclear fusion also have a potential spinoff in the biotech field. They are significantly more powerful than the lower temperature superconducting magnets currently used in MRI machines. Scientists at the National High Magnetic Field Laboratory in Florida have successfully tested an HTS magnet wound with REBCO that generated a magnetic field of 32 tesla (compared to 2–3 tesla for the magnets in an MRI machine).[38] MRI machines built with REBCO magnets would not require liquid helium as a coolant, and they could use smaller magnets (Mukoyama *et al.*, 2018). The market for MRI imaging devices is expected to reach $8B by 2027.

4.3.8 Life sciences and medical treatment

TAE Technologies is an example of a fusion company using its technology commercially in life sciences. Using advances in neutral beam technology, TAE was able to convert its technology into a cancer treatment known as boron neutron capture therapy. This treatment involves injecting non-radioactive boron into a cancer patient's tumor, then directing a neutron beam at the tumor. As with PBRT, healthy tissue is minimally damaged in the process. TAE Technologies' subsidiary, TAE Life Sciences, has raised \$70M since 2018 for this venture.[39]

5 A Fusion Energy Megafund

To address the challenges of fundraising for developing therapeutics in the biomedical industry, Fernandez *et al.* (2012) developed the concept of a "megafund" that uses a unique securitization structure to pool together the risk of pharmaceutical assets to offer investors a wider range of potential investments in this sector. This structure creates several tranches that provide investors with debt or equity through "research-backed obligations," RBOs, which pay out through a priority rule of cash flow distributions. The amalgamation of assets into a single structure promotes diversification and hedges against the individual risk of single assets, while the different characteristics of its tranches expand the investor cohort beyond the traditional reach of early-stage venture capital to include fixed-income investors, who control a much larger pool of capital. For example, in 2018, \$254B was invested globally through venture capital financing, while the size of the global bond market stood at \$102.8T, more than 400 times larger. The ability to finance through debt also has the advantage of a lower cost of capital than equity due to the tax deduction on interest expenses.

Fernandez *et al.* (2012) demonstrated that this novel financing method is able to produce quality returns for investors across both its equity and debt tranches in the oncology space. Fagnan *et al.* (2014) expanded this analysis into orphan drug development, which has higher probabilities of success and shorter time horizons, where it exhibited double-digit expected rates of return. More recently, Lo and Siah (2021) recognized the necessity of exploring correlations in transition probabilities, modeled through a single-factor model with a Gaussian copula, and found still promising (albeit less attractive) returns for an RBO structure with three tranches (senior debt, junior debt, and equity).

Generalizing the megafund approach to encourage investment in other necessary sectors that may struggle to receive sufficient capitalization, Hull *et al.* (2019) provide an extrapolated technique for securitizing high-risk "long-shot" projects, such as developing treatments for cancer and Alzheimer's disease and technologies that can address global warming. They define the criteria for a venture to be characterized as a long-shot asset as follows: "(1) it requires a large amount of initial capital; (2) it has a long gestation lag, during which no cash flows are generated and/or additional capital investment is required; (3) it has a low probability of success; and (4) if successful, its payoff is very large relative to the initial investment" (Hull *et al.*, 2019, p. 10). Nuclear fusion startups thoroughly adhere to this definition. The development of a power plant is expensive, the timeline to net positive energy output may take 10 to 20 plus years, and the technical complexity means each individual venture has a relatively low success rate, but the potential financial upside for an exothermic reactor is astronomical in an expected multi-trillion-dollar clean energy market.

In the first application of the megafund approach outside of the pharmaceutical sector, it is significant to consider and incorporate the five characteristic differences between investing in drug development and fusion ventures. First, the makeup of assets in a fusion megafund is more diverse than in a biomedical megafund, and thus requires a greater granularity in modeling. Pharmaceutical trials that treat a single disease in modeled form tend to converge more closely to independent and identically distributed (or correlated) random variables, while the intrinsic variety in fusion assets, because of their differences in technological approach, terminal business model, and even value proposition, make a similar approach illogical. We overcome this hurdle by classifying companies into different asset classes that exhibit roughly similar features.

Second, the conditional payout on terminal success is more certain for drug trials than for fusion ventures. While the pricing of biomedical IP has often been studied, and occasionally even regulated by governments, the valuation of a successful fusion company can only be estimated through comparable energy sources, since no fusion company has reached the implementation stage, and few transactions exist to date. Additionally, valuations will still depend on the market penetration of the specific fusion asset in question (e.g., directly competing in the energy space compared to licensing power plant designs).

Third, as the pharmaceutical industry is a legacy industry, there exists a rich dataset associated with drug trials. In comparison, with limited information on the outcomes of private nuclear fusion companies, the projection of future probabilities of success in this space is more difficult. The heterogeneity in fusion asset characteristics also calls for more individually specific metrics.

Fourth, a large proportion of fusion research globally is carried out at universities, or at subdivisions of larger research labs, while the pharmaceutical industry has an established commercial infrastructure for drug development. Thus, our model will need to incorporate a translation vehicle that can execute and finance the creation of a private legal entity that spins out of laboratory research.

Finally, and perhaps most significantly, while the value of drug development is typically limited to the outcome of the treatment, breakthroughs made by fusion companies can have practical applications in other fields and markets. This motivates the inclusion of spinoff valuation when evaluating fusion assets, and an incubation engine that can commercialize these indirect applications of core fusion research.

These five distinctive features are additional factors that must be incorporated into the securitization of fusion RBOs, the structuring of a fusion energy megafund, and the modeling of fusion assets. In this section we provide a detailed overview of how such a fund might be structured, and we relegate more technical details to the Appendix.

5.1 Modeling fusion assets

Fusion projects can be broken down into sequential phases. While the specifics of the stages will depend on the startup's chosen business model (discussed in Section 5.1.1), a common phase trajectory would have the first phase involve fundamental research and the development of IP, the second phase the development of an experimental prototype, the third phase the development of a scale device for testing purposes, and the final stage the deployment and operation of a fusion power plant. One phase cannot be launched until the successful execution of the previous one, and any promised cash flows will start only upon the success of the previous phases. We also assign a probability to each phase that represents the likelihood the company will be able to progress to the next stage. The beginning of each phase also corresponds to an additional round of financing that will cover all the company's costs for that period.

While we recognize that in practice many startups will need to seek bridge financing if they have underestimated their budget or their next promised objectives are delayed, we take the cost of each stage to reflect the cumulative financing for that phase. Figure 5 exemplifies a four-phase trajectory for a fusion asset by displaying the period, transition probability, cost/raised capital, and cash flow. To compute the net present value

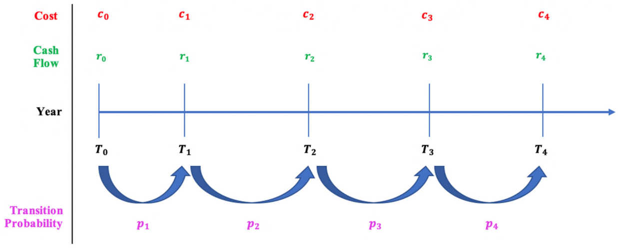

Figure 5 Overview of a four-stage trajectory of a fusion asset.

of the asset, we can simply take a discounted weighting of future cash flows from the present state.

Because many fusion companies are researching related technologies and employing similar business models, we introduce asset correlation into our model. We first assign a baseline initial correlation, since all the assets are in the fusion space. We then consider a vector of each asset's design attributes based on the decision tree in Appendix A (Figure A.2). To determine the pairwise correlation, we simply take the fraction of attributes the two assets share over the total number of attributes and add that number scaled to the baseline correlation.

Regarding cash flow, it is fair to assume that fusion startups will not generate any fusion energy cash flow until the terminal stage, when they are either producing power, selling reactors to energy producers, or licensing IP. However, spinoff applications of the core technology will allow for cash flows in the near term. At a cursory glance, the simplest means to determine cash flows for spinoff applications is to calculate the cost to license the IP rights to a company in the spinoff industry. Thus, the problem of predicting cash flows reduces to an accurate IP valuation. We experimented with several different IP valuation methods—among them, competitive advantage, premium pricing options, incremental cost savings, and market value multiples—in consultation with MIT's Technology Licensing Office, but we found vastly different results based on the methodology used. Due to the lack of consensus on the ideal valuation method, we abandoned this approach for a different procedure.

In consultation with fusion researchers, we determined a measure of relatedness for each spinoff direction from the fusion asset on a scale from 0 to 1. For each promising spinoff direction above a certain threshold of relatedness, we applied the phase trajectory model to a hypothetical company in the spinoff application industry. These parameters are based on traits of the industry, meaning the phase lengths will typically be shorter than for fusion, the costs smaller, and the cash flows more immediate. We can then determine the expected cash flows from these spinoff applications by weighting the computed cash flows of each spinoff with its measure of relatedness.

5.1.1 Asset classes

To account for the lack of homogeneity among fusion startups, we partition them into several asset classes that possess similar investment characteristics. The objective of this classification is to apply parallel phase trajectories to all assets within the same class, since they should require similar financing, bear similar risk profiles, and offer similar payouts.

Asset Class 1 contains companies that are experimenting with fusion designs and intend to produce fusion power plants. We further divide these assets into subclasses based on the asset's expected long-term principal source of revenue. **Asset Class 1A** contains companies intending to build reactors and power plants to produce energy directly to an electrical grid. **Asset Class 1B** contains companies intending to design and build power plants to be sold to existing energy companies or governments. **Asset Class 1C** contains companies intending to invest and license IP and academic discoveries.

Asset Class 2 contains companies that produce complementary and auxiliary technology for fusion power plants that is necessary for the successful deployment of a power plant. Examples of assets in Class 2 were described in Section 4, and include (but are not limited to) lithium-containing blankets, materials, control systems, fuel (e.g., deuterium and tritium in primary use cases), 3D printing technology, and cleanup technology. **Asset Class 3** contains companies producing nuclear fusion-propelled rocket engines. Finally, **Asset Class 4** contains university laboratories researching nuclear fusion or complementary areas that produce work with potential for commercialization.

5.2 *Holding company structure*

The proposed fund will be structured as a holding company that invests and owns interest in several subsidiaries, each representing an entity from one of the asset classes. Put simply, the holding company will raise money from investors and philanthropic partners through specialized financial contracts, acquire a series of assets—for example, ownership in a startup, licenses for IP, and certain rights—in exchange for capital from fusion-related entities, and repay investors with cash flows generated from the assets based on the terms of the financial contracts. The holding company will also have certain managerial roles and responsibilities. In addition to acquiring assets, it can sell or trade existing assets, it must aid in the operation of existing assets, and if certain assets are actionable rights, it can choose to execute such rights. The financial benefit of the holding company structure is that it allows managers the ability to facilitate debt and equity financing without the need for a separate special-purpose vehicle. Figure 6 depicts the generalized form of a fusion megafund. We qualify that, while the structure appears complex, the diagram simply aims to enumerate all potential asset types and investable vehicles for completeness. A fusion megafund in practice may not contain assets in certain categories, and it may include simpler financial relationships with startups.

As a means to increase the range of options for potential investors regarding risk and expected return, the holding company will offer financial contracts with differing terms that are repaid by a priority rule. This allows for both debt and equity investing, and a choice of risk with debt contracts

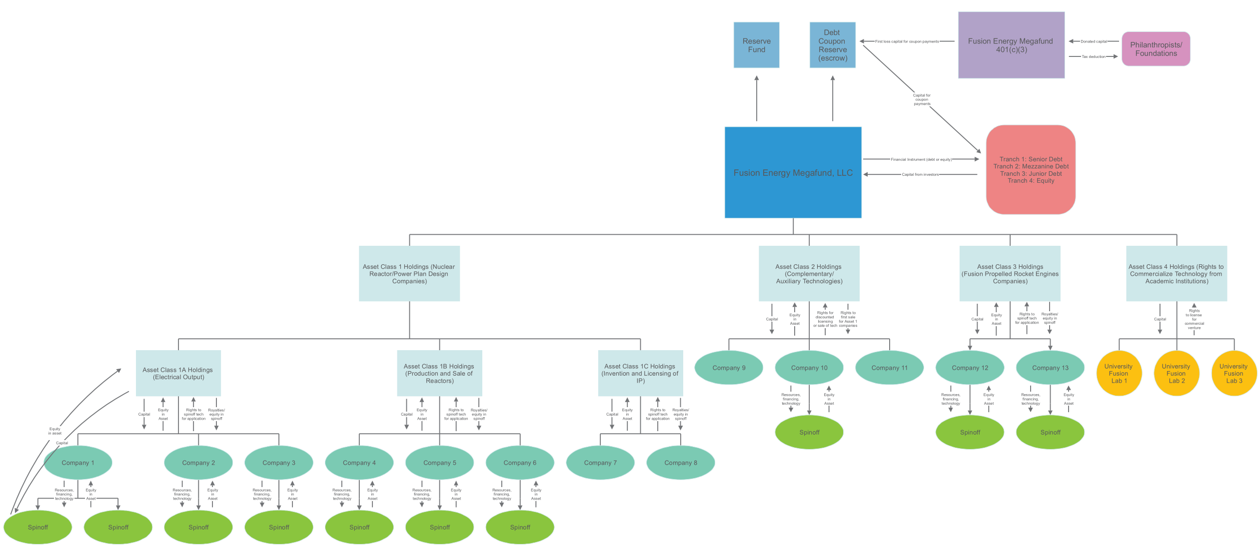

Figure 6 Proposed holding company structure for a fusion energy megafund.

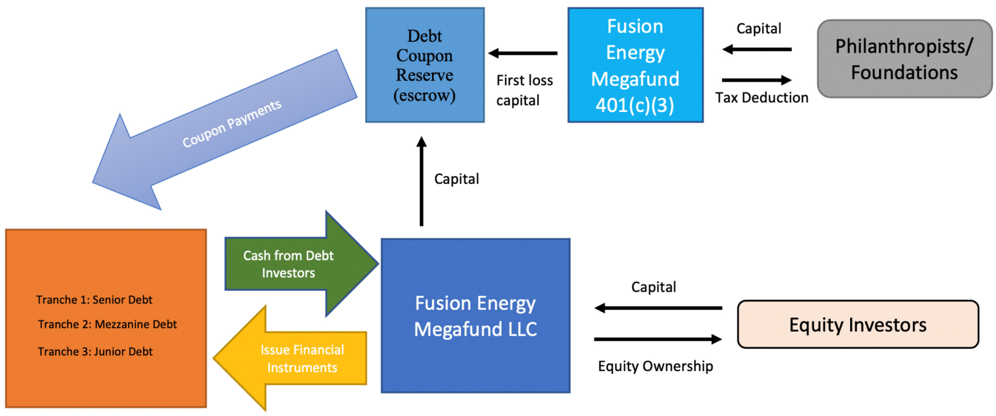

Figure 7 Illustration of a securitized fundraising model, with three debt tranches, an equity tranche, a philanthropic first loss capital mechanism through a 501(c)(3), and a debt coupon reserve.

based on the investment tranche. The holding company could theoretically have any number of tranches, but we consider a structure as illustrated in Figure 7 that has three debt tranches—senior debt, mezzanine debt, and junior debt—and an equity tranche. All debt tranches are promised a regular coupon return (typically annually) and a repayment of the investment at maturity. Based on the priority rule, the cash flow from assets will first be distributed to the senior tranche's coupon, then to the mezzanine tranche's coupon, and finally to the junior debt's coupon. In periods where there is a surplus of cash flow from assets, the priority rule will not apply, but if insufficient cash flows are generated to compensate the bondholders, then the more junior tranches will be at a greater risk of capital loss. As a result, the more senior tranches are promised a lower yield with lower risk, and the junior tranches a higher yield with higher risk. Credit assessment, yields, and interest rates will be determined by credit rating agencies such as Moody's, S&P, and Fitch. As an illustration, an example of a three-tranche bond structure might be:

- Tranche 1 (D_1): 4% yield, Senior Debt (AAA).
- Tranche 2 (D_2): 6% yield, Mezzanine Debt (BBB).
- Tranche 3 (D_3): 10% yield, Junior Debt (B).

The equity component will be designed as a private illiquid investment. This form of investment appeals to venture capitalists and institutional investors who have longer-term return objectives. Equity investors will be locked in for a period of more than 10 years, and the fund's life can be extended for another period, if needed. Unlike debt investments, equity investments have no promised payouts, but rather their investors receive all the remaining cash flows once the bondholders have been completely compensated. While it is the riskiest tranche, the equity tranche also possesses the highest upside, and could theoretically generate a colossal return.

We also include a mechanism for philanthropic support in our model, in which the donor provides capital to a separate 501(c)(3) with no expectation of financial return other than the tax deductibility on the contribution. Such donations will engender a philanthropically derived tranche of "first loss" or "catalytic" capital to compensate investors for risk and the relatively long duration of returns on fusion investments. This functions as a guarantee on coupon payments to senior tranche bondholders over the first couple of years. We estimate the philanthropic first loss component would require initial funding of at least \$500M to guarantee 5+ years of payments to senior tranche investors.

To further address the likely lack of early cash flows, even with successful spinoffs, we supplement the philanthropic first loss capital with a certain amount of initial investor capital into a debt coupon reserve, which is held in escrow. If the cash flows from assets and the concurrent philanthropic first loss capital fund cannot cover the coupon payments for the first few years, the capital held in the reserve is automatically distributed to bondholders to prevent default. Additional safeguards could also be added to the financing structure. For example, a financial guarantor, such as an insurance company, may be hired to reimburse investors in case of insolvency for an exchange of premium payments.

See Appendix B for additional details on modeling fusion assets and the holding company structure.

5.3 Megafund simulations

To assess the theoretical performance of the holding company structure, we generate a megafund portfolio of fusion companies, and use Monte Carlo methods to simulate its performance. A sample of 18 real-world companies[40] is used to construct a portfolio by acquiring a 10% share in each asset. The tenor of the fund is 20 years, chosen to reflect the long horizon for fusion companies to reach commercialization. Whenever capital is raised via debt, we issue three tranches: a junior tranche with a 6% interest rate, a mezzanine tranche with a 4% interest rate, and a senior tranche with a 2% interest rate, along with equity. We assume the total capital raised is $3 billion, using three different combinations of equity and debt, in the percentage allocations of 100:0, 60:40, and 40:60.

The parameters required by our simulation are the probabilities of success for phase transition, the length of each phase, the cost of each phase, the valuation of the asset at each stage (i.e., the cash that could be generated by selling it), the probabilities a company can generate a spinoff in a specific domain, and the correlations between projects. The four phases of our model consist of R&D, prototyping, breakeven energy generation, and commercialization. Surveying both plasma physicists and energy investors, we compiled parameter sets for each of the 18 companies, deriving a conservative estimate and an aggressive estimate from their expert opinion. Additional details on the simulation methodology and its parameters can be found in Appendix C.

Table 1 presents the performance statistics of the megafund. Under conservative assumptions, we find that the annualized return on equity (ROE) of the megafund has a range between −0.78% and 2%, and that its annualized Sharpe ratio has a range between −0.02 and 0.11. When a larger percentage of capital is raised via debt, there are more scenarios in which the megafund defaults. For an equity-to-debt ratio of 60:40, we find there is a low probability of default for the junior tranche (0.01%) and no default scenarios for the mezzanine or senior tranches. For larger amounts of capital raised using debt, however, we observe higher default probabilities (38.12% for junior debt at a 40:60 equity-to-debt ratio), because less cash is available for coupon payments. At the end of the megafund's tenor, the number of successful projects is small due to the lengthy duration of the projects: four projects on average are in the breakeven phase and three projects are in the commercialization stage.

In comparison, under aggressive assumptions, the annualized ROE ranges between 7% and 29%, and the annualized Sharpe ratio ranges between 0.20 and 0.24. The shorter lengths of each phase lead to three companies, on average, reaching the commercialization phase, and one finishing all stages. In addition to a much higher ROE from

Table 1 Performance statistics of the megafund under (a) conservative and (b) aggressive parameters. The ratio of equity-to-debt is varied from 100:0, 60:40, and 40:60.

Measures	100:0 (Equity:Debt)	60:40 (Equity:Debt)	40:60 (Equity:Debt)
(a) *Conservative parameters*			
ROE (total)	39.12	11.85	−14.50
ROE (arithmetic avg)	1.96	0.59	−0.73
ROE (geometric avg)	1.66	0.56	−0.78
Pr(ROE $< -10\%$)	3.36	81.16	87.79
Sharpe (annual)	0.11	0.02	−0.02
# Prototyping	0.0	0.0	0.1
# Breakeven	4.2	4.3	4.3
# Commercialization	3.4	3.1	2.7
# Finished	0.1	0.1	0.1
Junior default (%)	0.00	0.01	38.12
Mezzanine default (%)	0.00	0.00	2.79
Senior default (%)	0.00	0.00	0.01
(b) *Aggressive parameters*			
ROE (total)	283.17	415.12	584.44
ROE (arithmetic avg)	14.16	20.76	29.22
ROE (geometric avg)	6.95	8.54	10.09
Pr(ROE $< -10\%$)	5.45	21.45	28.53
Sharpe (annual)	0.24	0.21	0.20
# Prototyping	0.0	0.0	0.0
# Breakeven	0.1	0.2	0.3
# Commercialization	3.1	2.8	2.6
# Finished	1.2	1.2	1.2
Junior default (%)	0.00	2.43	14.56
Mezzanine default (%)	0.00	0.87	6.62
Senior default (%)	0.00	0.24	2.04

the post-commercialization valuation, the shorter phase lengths under aggressive assumptions add greater certainty to the project outcome, since more companies will reach the riskiest final transition. However, because more projects will also fail commercialization under these assumptions (in addition to more projects commercializing successfully), the probability of default for all tranches is higher at a 60:40 equity-to-debt ratio. This also holds for the more senior debt in a portfolio with a 40:60 equity-to-debt ratio.

In our model, when a project successfully finishes the R&D stage, there is a potential for spinoff companies to use the new technology developed during this phase. We give the number of successful spinoff projects modeled in our portfolio in Table 2. We find that the final portfolio contains, on average, four spinoffs specializing in laser technology, five in medical scanning, and six in magnet technology. The megafund portfolio performance thus demonstrates visible dependence on cash flows from these licensing deals.

Table 2 Average number of successful spinoff companies from the megafund portfolio.

Spinoff Companies	Average Number	Standard Deviation
Laser	4.0	1.8
Magnets	6.0	1.0
HTS	0.9	0.3
Cyclotron	3.0	1.0
EM aircraft launch	2.3	1.6
Thermal measurement	2.3	1.6
Robotics	1.7	1.8
Medical scanning	5.5	1.9
Life science	1.7	0.8
Materials science	1.7	0.8

Appendix C includes additional results that explore the sensitivity of the megafund's performance to valuations of projects post-commercialization, the cash flows from spinoffs, the phase transition probabilities, the costs associated with each phase, and the lengths of each phase.

6 Discussion

Beyond the increased availability and access to capital, a fusion megafund will serve the fusion sector by aiding in the development of a modular fusion ecosystem, improving the chances of commercial viability and promoting greater technological exploration. Many fusion startups require the use of auxiliary technologies, but they lack the time and resources to develop optimized components. For example, all projects using tokamaks need software to predict and address plasma disruptions, which could be developed most efficiently by a third party. Similarly, while some private sector development has begun on lithium cooling blankets, improving the quality of the product would have a widespread impact on all fusion companies. Given the already immense technical challenges that fusion companies face in generating breakeven energy, developing a modularized supply chain for auxiliary components is necessary for any eventual fusion ecosystem. Additionally, by promoting translation from university research, the megafund will be able to test the commercial viability of more innovative and efficient approaches. Early-stage concepts that may struggle to receive interest from investors will also receive funding, as the megafund will seek to diversify both the stage and risk of all assets under management.

For investors, the most evident value of the megafund is de-risking fusion investment through large-scale diversification, providing them with a means to invest in the entire sector, including auxiliary technologies. Since there is no consensus on the ideal approach for a fusion company to achieve success, investing in only a single concept is inherently risky because of the uncertainty regarding which technology ultimately will succeed. By pooling a large number of high-risk fusion investments into a single portfolio, the investment risk associated with any single fusion project can be significantly reduced. In the simplest terms, the megafund transforms an investment from a bet on an individual design to a bet on fusion energy. The holding company structure also alleviates the burden on individual investors to expend substantial resources to obtain the highly technical knowledge required to make a prudent investment in the fusion space. As noted, the investor community is largely unfamiliar with the scientific aspects of fusion and has difficulty discerning the merits of different technological approaches. Because the megafund will have deep-rooted knowledge of the fusion space, it will select a strong cohort of fusion companies in which to invest, and institutional investors, high-net-worth individuals, and sovereign wealth funds will no longer be required to possess sector-specific expertise.

Second, through securitization, the megafund will appeal to a wider range of investors. The current mechanism of funding fusion almost entirely through equity investments in individual companies requires a high degree of risk tolerance and a long time frame for any return on investment. The megafund structure provides debt tranches with a lower cost of capital than equity, which greatly increases the likelihood of generating a return, especially with the addition of a first loss capital fund and a debt coupon reserve fund. The other consideration of investors, the ability to cash out their investments in a reasonable time horizon, is also addressed. The securities can be arranged to have varying maturities, permitting return horizons of 3 to 5 years, as long as a sufficient number of equity investors exist. Given the high costs of development, the fusion industry will otherwise struggle to grow when the majority of investment funds participating have \$100M to \$500M in AUM, which is around the average cost to build a test reactor.

The fusion megafund was specifically designed to appeal to large sovereign wealth funds, which can place a long-term macro-bet on the fusion industry as a hedge against oil depletion or stricter government regulation enforcing the increased use of renewable energy. In fact, five of the largest sovereign wealth funds are long oil and gas: the Government Pension Fund Global of Norway (AUM: \$1.2 trillion), the Abu Dhabi Investment Authority of the United Arab Emirates (AUM: \$580B), the Kuwait Investment Authority (AUM: \$535B), the Public Investment Fund of Saudi Arabia (AUM: \$380B + 98.5% ownership in Aramco), and the Qatar Investment Authority (AUM: \$335B). Recent transactions among these funds have been trending toward investing in renewable energy. For example, Norway's Parliament voted in June 2019 to sell \$13B worth of assets from oil, gas, and coal extraction companies, and move up to \$20B into renewable energy companies. In 2018, the Public Investment Fund and Sanabil Investments became a 25% shareholder in the largest renewable energy producer in the Middle East, ACWA Power, and is planning to increase that share of ownership to 40%. In June 2020, Nordic Pension Funds agreed to invest \$1.5B in what is set to be the largest renewable energy portfolio in the world. Nuclear fusion warrants a place next to solar and wind in any of these renewable energy portfolios.

Third, the megafund also allows for robust participation from the government and philanthropists in an effort to further de-risk private sector investment. As demonstrated with nuclear fission, governments have historically guaranteed or insured riskier projects in the energy sector. Through the megafund, governments can provide directed capital and resources, or a guarantee of first loss position. Furthermore, from a macro-perspective, governments looking to establish an actionable fusion program that develops pilot plants (as is being discussed in the US) can facilitate the initiative through a public–private partnership involving either the formation of a national fusion energy megafund or a services contract with an existing fusion energy megafund.

Regarding philanthropy, first loss capital guarantees have been effective mechanisms for encouraging private investment in initiatives operating in unproven markets (often in low-income countries) and in projects with uncertain revenue streams. The philanthropic community is already awash in funds targeting climate change. In one recent example, Jeff Bezos created a \$10B Earth Fund, led by former World Resources Institute CEO Andrew Steer, which will invest in scientists, non-governmental organizations, and the private sector to further the development of new technologies. Growth in foundation AUM and limited options for the deployment of large-scale philanthropic capital suggests that a well-articulated fundraising effort for fusion would find broad appeal.

Fourth, by allocating spinoff rights to investors, the megafund would be able to separate the financial interest of investors from the social and altruistic goals of researchers. While many researchers employed at fusion startups would not want to divert their attention from energy research to external applications, investors would be delighted to effectuate the licensing of technology for a profit while still supporting energy research.

Lastly, through its treatment of auxiliary technologies, the megafund offers investors the opportunity to invest in lower risk assets that exploit niche areas of the fusion space and provide the capital necessary to monetize the products and services needed in a fusion ecosystem. Private fusion does not have to be an unmitigated success for auxiliary technologies to produce significant financial returns for investors. All government-sponsored fusion initiatives, including ITER, would benefit from the incorporation of these products and services, and could serve as immediate customers.

However, there are several challenges and limitations to the megafund approach. The principal challenges in deploying a fusion energy megafund are in its operation and execution. Its operational impediments include sourcing a large amount of initial capital from investors (our estimates are in the $10B range), managing the complexity of a highly dynamic financial structure, and successfully carrying out the processes of both incubation and translation.

In regard to fundraising, the large sum raised by ITER's government sources is evidence that large-scale capitalization is indeed feasible in this space. Given the annual investment in the overall energy sector, it is quite reasonable to expect that the private sector could spearhead an initiative at less than half the cost of ITER. In addition, both nuclear fission and the Human Genome Project saw large infusions of private capital once research was proven in public projects.

In managing the megafund's day-to-day operation, we acknowledge the challenges of investing in competing and highly sophisticated technologies within the same industry. However, EIG Global Energy Partners, an institutional investment firm specializing in the energy sector, already has $22.4B under management for 363 portfolio companies across 36 countries, employing 58 investment managers.[41] The scope of the fusion energy megafund, in terms of its subsidiary companies and international presence, would be much more limited.

Finally, while incubating a spinoff certainly requires a robust infrastructure and an abundance of resources, the process of generating divestitures is well documented in the business community, and it has already been displayed in the fusion space through the example of TAE Life Sciences. Translation as well has been one of the primary means for the genesis of fusion companies, and the development process of Commonwealth Fusion Systems, for instance, could be used as a paradigm.

The obstacles to the proper execution of a megafund, on the other hand, are slightly more trying. These include the limited quantity of active fusion ventures, the difficulty in estimating the future value of a power-producing fusion plant, and the dearth of data to derive parameters for accurate financial simulations.

Even though the fusion industry at present is rather small, recent trends have shown an annual increase in new fusion initiatives. The inclusion of auxiliary technologies and an effective translation vehicle for university laboratories would serve as a mechanism to increase the number of active companies operating in fusion without oversaturating the space.

To address the future value of a successful fusion company, we built a compound real options model that calculates the potential value of fusion power plants. This method, however, proved very sensitive to input variables, much like the results of Bednyagin and Gnansounou (2011). They compared the estimated valuation of ITER's research and development efforts using a basic Black–Scholes/Merton real options model, a compound real options model, and a fuzzy real options model. The results were "in the range between $103 billion and $292 billion with the least possible downside value of –$36 billion and the least possible upside value of $503 billion."

As an experiment, we used a Black–Scholes/Merton real options model to estimate the value of Commonwealth Fusion Systems' ARC tokamak, with preliminary results in the range of $120B. However, sensitivity analysis showed that the length of the project phases, the future evolution of electricity prices from 2040 onward, the expansion rate of the sector (i.e., the number of power plants built per year), and the operation and maintenance costs all significantly influenced the option value. Moreover, the track record of fusion projects in underestimating the necessary time and budget highlights the low reliability of input data in the eyes of potential investors. As a result, we concluded that the option valuation method was too volatile to be useful.

We encountered similar problems in determining the transition probabilities for fusion assets. While studies advocating for megafunds in the biotech sector are able to generate sample portfolios and credibly simulate their returns, the relative infancy of the fusion industry makes it very hard to obtain analogous parameter values, and the diversity of fusion assets means that the assumption of independent and identically distributed random variables in an unmodified simulation is erroneous. Small variations in the estimated transition probabilities would have large influences on the present return. We attempted to associate historical triple products to predict both future triple products and company reinvestment, but we found little correlation given the sparsity of the data.

We overcame these difficulties by collaborating with a number of stakeholders in the fusion industry to define a practical implementation of the megafund. We avoided the challenges of generic parameter estimation by constructing various simulations of the fund around real fusion assets, with attribution and predictive data generated through expert opinion, rather than using a set of generalized assets.

To summarize the performance of our simulation, we find that there are positive returns on equity and low default rates for the capital raised using debt. However, these are highly dependent on the value of successful and in-pipeline projects, and the cash flows from licensing deals of potential spinoff companies. Additionally, the megafund's performance is correlated with assumptions about the length of time to project commercialization and the transition probabilities between project phases. Under conservative assumptions, the 100:0 equity-to-debt portfolio's annual average ROE was 1.96%, while under aggressive assumptions, its annual ROE was 14.16%. Our sensitivity analysis presents a range of parameters as a benchmark to associate our assumptions about the private fusion sector with predicted returns under the securitized holding company structure. Nonetheless, the positive ROE under most parameters demonstrates the success of the megafund structure to financially engineer success in an asset class that has historically been characterized as possessing a "valley of death."

7 Conclusion

Now that the NIF at Lawrence Livermore National Laboratory has achieved scientific net energy and many fusion startups are poised to do the same by the end of the decade, it is the ideal time for investors interested in the fusion space to act. As more of these major milestones in fusion are achieved, the valuations of fusion companies and the demand for their auxiliary technologies will almost certainly increase. While the megafund model for the fusion industry still possesses some imperfections, it nonetheless addresses investor hesitations about risk, return horizon, and asset class. Ultimately, the financial incentive of investment in fusion should be tempered by principled motivation. Of the sustainable energy sources necessary to bring about a decarbonized energy system, it is fusion energy that possesses the highest financial upside. While the prospect of raising billions in funding for a still unproven industry may seem too costly and unrealistic at this time, the cost to society of not innovating around sustainable energy may prove even more expensive.

Appendix A A Primer on Fusion Energy in Theory and Practice

A fusion reaction is one in which two atomic nuclei amalgamate, or "fuse," into the nuclei of new elements (see Figure A.1), along with subatomic neutrons and protons. There are hundreds of such reactions, but in those of interest to fusion energy production, the reaction releases large amounts of net energy, due to the mass of the final products being smaller than the total mass of the original particles. This is a consequence of Einstein's famous discovery of mass-energy equivalence, $E = mc^2$. The net energy is manifest in the kinetic energy of the products.

Fusion energy lacks many of the drawbacks of other energy sources. In contrast to fossil fuels,

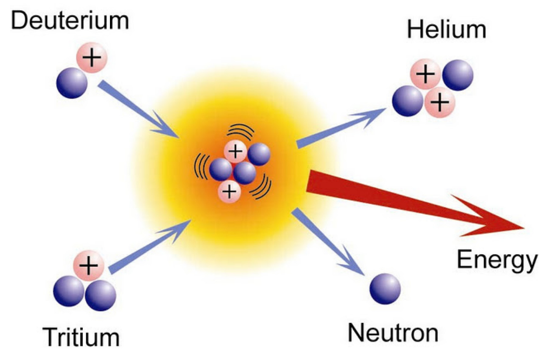

Figure A.1 Illustration of a fusion reaction with deuterium and tritium as the reactants.
Source: chemwiki.ucdavis.edu.

its reactants are virtually unlimited and non-polluting, and it has a smaller environmental impact. In contrast to renewables, it is scalable and it has a higher power density. In contrast to nuclear fission, it produces a lower level of waste, has no chance of a meltdown, and does not require the use of fissile materials, which must be controlled due to concerns over nuclear proliferation. Furthermore, the predominant fuels used in fusion reactions (deuterium and lithium, which can be used to produce the other reactant, tritium) are almost entirely innocuous and are readily available. The stable hydrogen isotope deuterium poses no health threat, and its supply is virtually unlimited, as it can be separated from Earth's oceans. Tritium is a naturally occurring radioactive isotope of hydrogen that can be produced internally within a fusion power plant, using the relatively common element lithium distributed around the fusion-producing volume. Despite its radioactive properties, humans are already exposed to trace amounts of tritium daily, through the food and water they consume (the average American receives about 310 millirem annually from such natural sources).[42] In addition to availability, fusion power plants require only a small quantity of fuel: a 500-MW fusion plant would only require, on average, 50 g of deuterium per day, and some magnetic fusion devices would only need 0.2 g of fuel in the core plasma.

These features give fusion energy a significant advantage in terms of levelized cost of energy, once overnight cost targets are attained (IRENA, 2017).

However, the technical challenges to develop commercial fusion plants are significant. Scientists and engineers worldwide have been working for decades to demonstrate the viability of this technology. As characterized by the World Nuclear Association, "fusion power offers the prospect of an almost inexhaustible source of energy for future generations, but it also presents so far insurmountable engineering challenges."[43]

A.1 History of fusion research

The study of fusion originated with English scientist Sir Arthur Stanley Eddington, who suggested that stars produce energy through the combination of hydrogen atoms into helium, a process later known as the proton–proton chain. In 1933, the first fusion reaction was observed in a laboratory setting when the New Zealand physicist Ernest Rutherford's student, Mark Oliphant, used an accelerator to fire deuterium toward a target, revealing the nuclei of helium-3 and the radioactive hydrogen isotope tritium. In 1939, the physicist Hans Bethe was awarded the Nobel Prize for his astrophysical work demonstrating two fusion reactions that occur in stars, the carbon–nitrogen cycle and the proton–proton cycle.

From the late 1940s into the 1950s, scientists began exploring whether they could replicate these astronomical processes for commercial energy production. The United Kingdom Atomic Energy Authority filed the first patent for a fusion reactor employing the Z-pinch technique, a method in which a magnetic field generated from an electrical current would compress the plasma where a fusion reaction takes place. Soviet scientists Igor Tamm and Andrei Sakharov simultaneously proposed another magnetic confinement device, the famed tokamak, which used a plasma containment chamber in the shape of a torus. Another major magnetic confinement design, the stellarator—proposed by Lyman Spitzer in 1951—restricts plasma through a series of wrapped magnetic coils. Inertial confinement, in which lasers compress a fuel pellet to initiate fusion, was first proposed at the Lawrence Livermore National Laboratory in 1965, soon following the invention of the laser.[44]

After the Soviet T-3 tokamak showed promising results in 1968, exhibiting a 10-fold better performance than any other concept, and catalyzed by the energy crisis of the 1970s, three major tokamak projects were initiated: the Tokamak Fusion Test Reactor (TFTR) in the US, the Joint European Torus (JET) in England, and the Japan Tokamak (JT-60) in Japan (Barbarino, 2020). At the same time, the International Atomic Energy Agency initiated increased collaboration between leading nations in energy research (chiefly the US, the USSR, and the UK) to form the International Tokamak Reactor Workshop in 1979. These efforts to galvanize leading nations to develop fusion reactors in a cohesive manner culminated in the genesis of ITER, first proposed in 1987 and launched in May 1988.

Many experimental tokamaks began publishing auspicious power outputs in the 1990s. In 1991, using deuterium–tritium fuel, the JET produced 2 MW of power. In 1994, the TFTR demonstrated 10 MW of peak power production. This was surpassed by a later experiment by the JET that resulted in 16 MW of peak fusion power, which still is the record for power release (Barbarino, 2020).

Recently, other experiments around the world have achieved major milestones, while novel fusion concepts are being explored due to an expanding private fusion industry (see Section

A.3). In 2016, the Alcator C-Mod compact tokamak achieved a new record in the fusion fuel pressure, exceeding 2 atmospheres for the first time.[45] In 2018, the Wendelstein 7-X in Germany demonstrated the highest triple product—the product of density, temperature, and confinement time—for a stellarator device, although still 10 times lower than the previously discussed tokamaks (Klinger *et al.*, 2019). In that same year, China's Experimental Advanced Superconducting Tokamak (EAST), using superconductor magnets, successfully demonstrated an electron temperature exceeding 100 million degrees Celsius, seven times hotter than the core of the Sun, for over 100 seconds.[46] In August 2021, the National Ignition Facility (NIF) at Lawrence Liverpool National Laboratory reported that its inertial confinement fusion facility generated 70% of the power required by the experiment, almost achieving ignition (Tollefson, 2021). Most recently, on December 5, 2022, the NIF achieved the first fusion reaction to produce net energy, making the theoretical possibility of fusion a reality.

ITER

Despite these advances, the attention of the fusion sector for the past two decades has been on ITER, which by the 2000s had become the world's largest fusion program, backed by the European Union, China, India, Japan, Russia, South Korea, and the United States. In 2005, it was agreed that the experimental device would be built in southern France, and construction was initiated in 2010. As of October 31, 2022, ITER claims to have completed 77.7% of construction. Current projections have ITER obtaining its first plasma in the late 2020s and beginning deuterium–tritium net energy operations in 2035.[47]

While ITER has deployed extensive resources to demonstrate its goal of a robust future fusion sector, the project has also been subject to censure due to logistical and engineering concerns. Many megaprojects have a recurring problem with underestimating their time and budget, an inevitable consequence of their intrinsic technical complexity. ITER is no exception. ITER has quadrupled its initial economic cost estimates, and to date has required an additional 10 years beyond its original schedule to deliver its first results. The US Department of Energy has publicly stated that it believes that ITER's presently stated budget (north of $24B) does not reflect an accurate estimate given the work remaining, and instead estimates it to be closer to $65B (Kramer, 2018). In 2020, the European Court of Auditors released a report foreseeing even further cost increases and delays, due to issues in its supply chain and challenges in its management.[48]

ITER's critics tend to focus on its large resource consumption and energy expenditure, despite not producing any electrical power as compensation. ITER's non-interruptible power drain has been estimated to be between 75 and 110 MW (Gascon *et al.*, 2012). The water demand to sustain ITER's fusion reaction is also extremely high, necessitating a flow rate of 180,000 gallons per minute (Jassby, 2020). Even with the capital and resources so far expended, ITER does not further the push toward economically viable fusion, with the cost for each megawatt of energy around two orders of magnitude greater than competitive energy sources (NASEM, 2021). While it should be noted that ITER was primarily designed for the scientific investigation of a fusion plasma that is predominately heated by its own fusion reactions, its direct economic prospects are not favorable.

Finally, many fusion researchers fear that the outcome of ITER's experiments will brand the fusion industry with an indelible reputation for requiring very large facilities and a long construction time, and generally poor economic prospects, thus potentially limiting future funding.

A.2 Technical challenges of fusion

Before analyzing the variety of reactor designs and technological approaches to fusion, it is essential to understand the technical challenges of achieving breakthrough energy. The core difficulties of fusion are all grounded in physics: heating plasma to over 100 million degrees Celsius, confining plasma at such a temperature, and sustaining the reaction with a sufficient combination of temperature, fuel density, and confinement. From these basic constraints, we can derive a number of secondary objectives required for the successful execution of a power plant, including (but not limited to) increasing the pressure of the confined plasma to fuse nuclei more rapidly and thus produce a sufficient fusion power density, discovering the means to handle the intense exhaust heat, and developing novel materials that can best withstand neutron-induced damage. Furthermore, while fusion is generally an "on-demand" power source, there are engineering complications regarding fusion power being generated for a finite duration, sometimes by design, due to the nature of plasma confinement, and sometimes unintentionally, due to plasma instabilities.

Fusion energy gain is symbolized by Q, the ratio between the power released by the fusion reaction and the power put into the fusing plasma. A Q value of 1 would signify that a reaction produces breakeven energy. The Q of a fusion plasma is determined by the Lawson criterion, which multiplies the density of fuel ions in the plasma and the energy confinement time plotted versus the temperature of the plasma fuel. Taken together, these three parameters are called the triple product. A commercially viable power plant would require a steady-state Q value above approximately 10. Before the exothermic reaction at Lawrence Livermore National Laboratory in December 2022, the world record was held by the JET tokamak mentioned earlier, which, through its production of 16 MW, demonstrated a Q of 0.6 (Jamieson and Highfield, 2009). ITER is planning on achieving a Q of 10 by generating 500 MW from 50 MW of input power. The SPARC tokamak under development by Commonwealth Fusion Systems plans to achieve a $Q > 2$, and an analysis of its mostly likely performance indicates a $Q \sim 11$.[49]

A.3 Diverse approaches to fusion energy

Beyond the scientific issues of generating and sustaining a fusion reaction, the difference of opinion about the ideal technological approach to a sustained fusion reaction has presented challenges for the fusion space. The principles of deuterium–tritium fusion energy production are well known, since they are set by the Lawson criterion, requiring a temperature of approximately 100 million degrees Celsius. However, since the density-confinement product can be achieved over a very large range of fuel density—a factor of a billion—this has naturally led to a contest in developing the optimal fusion approach. The two principal techniques under consideration are the use of magnetic fields, in magnetic confinement fusion, or the use of high-energy laser beams, in inertial confinement fusion. An additional category that combines aspects of magnetic and inertial confinement is known as magneto-inertial confinement.

Within the category of magnetic confinement, several designs exist for toroidal configurations. There are also open linear configurations, such as mirrors (Teller, 1981), which have recently garnered interest after a period of dormancy.[50] A tokamak is structured to have magnetic forces that spiral around the torus, thus creating a poloidal field, oriented along the "pole" of the torus, from a toroidal electric current, produced in the plasma. Tokamaks have by far the best plasma confinement performance (by a factor of at least 10 according to the Lawson criterion, compared to other concepts) yet they face challenges from instabilities and in producing a steady-state fusing plasma. There are estimated to be over 150 tokamaks around the world, most notably in large government projects (ITER, JET, JT-60) and several private startups (Commonwealth Fusion Systems, Tokamak Energy).

Other magnetic confinement designs have also been implemented. A stellarator uses a similar approach to a tokamak, except that the magnetic field is twisted exclusively by non-axisymmetric coils. Since there is no plasma current, operating the plasma in the steady state is more easily achieved, but unconfined particle orbits due to loss of toroidal symmetry can cause high transport losses (Xu, 2016). The most prominent stellarator implementation is the Wendelstein 7-X, and there are two recently founded private companies—Type One Energy and Renaissance Fusion—attempting to commercialize the concept. The pinch is a design that creates a magnetic field from the plasma current to compress (or "pinch") the plasma into a torus or cylinder (Liberman *et al.*, 1999). Pinches were used in some of the first attempts at fusion, and they are currently being pursued by the private venture firm Zap Energy. Two concepts for compact toroids include the spheromak and the field-reversed configuration, which confine plasma on closed magnetic field lines into a self-stable torus, possessing slight differences in the shape of the terminal. The company TAE has been leading the charge on experimental field-reversed configurations.

Inertial confinement also has several variations in design, as the implosion of fuel pellets can be accomplished by laser ablation via direct-drive inertial confinement, or by X-rays inside a "hohlraum," known as indirect-drive inertial confinement. The latter is epitomized at the NIF at Lawrence Livermore National Laboratory. Inertial confinement has been primarily funded for national security reasons, including at NIF. However, the Oxford spinout startup firm First Light Fusion has been using core principles from inertial confinement in its development program.

Magneto-inertial fusion synthesizes the two principal techniques by using both compressional heating and magnetization to reduce the thermal energy loss of the reaction. Private venture firms General Fusion and Helion Energy are known for developing designs based on magnetized targeted fusion, a sub-branch of magneto-inertial fusion.

While generally dismissed by physicists, cold fusion is a conjectured mechanism for achieving a fusion reaction that occurs at room temperature. In 1989, two scientists detailed the results of an experiment under laboratory conditions that produced unexplained amounts of excess heat, and byproducts of an apparent fusion reaction (neutrons and tritium; Ritter, 2016). Many studies have attempted to recreate the experiment or provide a theoretical justification of the results, but to no avail. Thus, cold fusion has been generally discredited apart from a small community of committed researchers, and subsequently receives very little funding. Google has commissioned several projects in recent years, but it has not discovered any corroborating evidence of cold fusion's existence. However, Google's experiments were valuable in their production of novel materials and tools such as calorimeters that can function under extreme conditions.[51]

The variation in fusion experiments and fusion startup companies goes beyond confinement techniques. Even within a given design category, the differences in execution can be large, as projects and companies often hone their efforts on solving one of the core challenges of fusion through the use of novel materials, equipment, or setup. With regard to fuel, while deuterium–tritium reactions

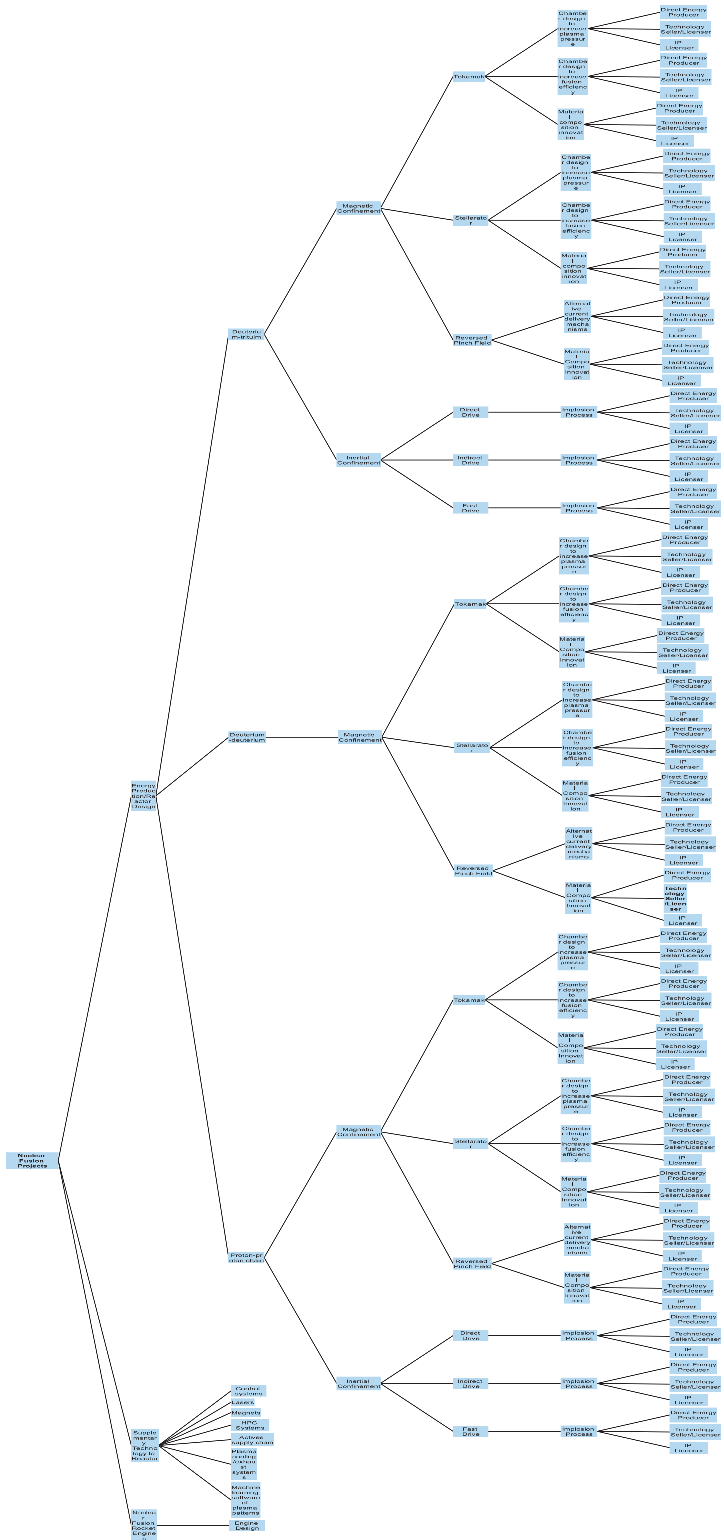

Figure A.2 A decision tree enumerating different technological and business permutations of a fusion-related startup.

(D-T fusion) may present the nearest term path to viable commercial fusion, deuterium–deuterium, proton–boron, and deuterium–helium-3 fusion are also being explored. To illustrate the differences among private fusion ventures in particular, we prepared a decision tree displayed in Figure A.2.

Since commercial fusion energy has yet to be demonstrated, there is still no consensus as to the ideal set of design attributes, and it is important to explore a number of permutations. At the present, D-T tokamaks offer the shortest expected time to achieve commercial fusion, based on the Lawson criterion and experimental results, and have thus been recommended as the leading concept for a US pilot plant (NASEM, 2021). Stellarators are seen as a leading alternative, based on achieved Lawson parameters. However, many of the aforementioned competing concepts, while possessing longer execution timelines and poorer confinement time results, aspire to better fusion performance and lower overnight and variable costs in the longer run. It is thus imperative to support a diverse array of concepts, with financial allocation based on the risk–return dynamics of each approach.

Appendix B Securitization

B.1 Asset phase progression

In our model, we apply discrete staged trajectories to all assets. Asset i is purchased at time T_0^i, and is subsequently evaluated at time $T_k^i, k \epsilon \{1, 2, \ldots\}$, where the time between T_{k-1}^i and T_k^i represents period k^i and is of length l_k^i. Note that the megafund begins at time $t = 0$ and progresses continuously. Since assets can be acquired at any discrete time step, it is not always the case that $t = 0 = T_0^i$. Additionally, because assets can have different progression cycles, it should also be noted that it is not invariant that $T_k^i = T_k^j$ for $i \neq j$, even if $T_0^i = T_0^j$. Time τ is the time for an asset (if it has not failed) to reach the independent operational threshold. This refers to a terminal stage such that the asset is able to fully and independently operate without the need for additional funding, and generates peak cash flows.

Each period has a Markov probability, p_k^i, that corresponds to the chance the asset reaches the next stage conditioned on the asset being in that state. For example, p_2^i represents the probability that an asset progresses from T_2^i to T_3^i, given the asset has already proceeded from T_0^i to T_1^i and from T_1^i to T_2^i. Thus, the probability that asset i progresses from infancy to stage k is equivalent to $\prod_{j=1}^{k} p_j^i$. After τ, we have $p_k^i = 1$.

Each stage k, up until the independent operational threshold is reached at time τ, also requires a corresponding investment of c_k^i. One can conceptualize these stages as venture capital investment rounds. Additionally, each stage has an associated cash flow, x_k^i. Note that, realistically, many early stages will have zero cash flows, or hypothetical realizations of intangible sub-assets, such as IP. After the independent operational threshold is reached, we model the future stages as having constant annual cash flows, X, or constant annual cash flows with polynomial drift, $X + \alpha(k - \tau)$. We can express the net present value of the asset by a discounted sum of current and future cash flows weighted by stage probability. For stage k, this net present value is equivalent to:

$$\sum_{j=0}^{k} x_j^i + \sum_{j=k+1}^{n} \left(x_j^i * \left(\prod_{m=k+1}^{j} p_m^i \right) * \gamma^{j-k} \right)$$

where γ is the discount factor or, equivalently, $1/(1 + r)$, where r is the temporal discount rate or cost of capital. Incorporating the independent operational threshold into the net present value computation, we get:

$$\sum_{j=0}^{k} x_j^i + \sum_{j=k+1}^{\tau} \left(x_j^i * \left(\prod_{m=k+1}^{j} p_m^i \right) * \gamma^{j-k} \right) + \sum_{j=\tau}^{n} ((X + \alpha(j - \tau)) * \gamma^{j-k})$$

For simple assets whose only value is derived from discounted cash flows in a terminal operational state, this reduces to:

$$\sum_{j=\tau}^{n} ((X + \alpha(j - \tau)) * \gamma^{j-k})$$

However, sources of more near-term cash flows are heavily relied upon in our analysis, including IP and spinoffs.

B.2 Spinoffs

We expect subsidiaries to create new independent companies, either divestitures or spinoffs, that use core technology developments to sell complementary or auxiliary technologies to generate cash flows for debt holders, and potentially equity holders. We denote by S the set of spinoff compatible technologies, where some elements may include superconducting magnets, MRI scanners, particle accelerators for cancer treatment, and thermal analyzers.

To project the performance of a spinoff from a particular company, we require a measure of the technological relatedness of the core technology to the spinoff, and a measure of the disruption of the technology in the spinoff market. We therefore define $\theta_{j,s}$, $j \epsilon \{1, \ldots, n\}$, $s \epsilon S$, to be a normalized rating from 0 to 1 of the applicability of company j's technology to spinoff market s. For example, a magnetic confinement fusion startup will have a very low θ for laser-oriented spinoffs (likely zero), but a high θ for magnetic-oriented spinoffs, like MRI scanners. The reverse would be true for inertial confinement companies.

In addition, we define $\varphi_{j,s}$, $j \epsilon \{1, \ldots, n\}$, $s \epsilon S$ to be an unrestricted, relative positive metric that is a high-level assessment of the ability of company j's technology to disrupt spinoff market s. This market penetration measure is a function only of the projected performance of company j's technology, and is unaffected by features of spinoff market s. The specific characteristics of the spinoff industry—barriers to entry, level of competition, composition of its customer base, and market capitalization—are incorporated into other analyses. Additionally, the quality of business operation of company j is unaccounted for in φ, since the spinoff may be subject to new leadership and governance under the proposed structure. Thus, φ is determined by comparing the primary value proposition—such as technological quality, novelty, or cost reduction—to an average startup in the space. A value $0 < \varphi < 1$ indicates that the fusion company's spinoff is below the average quality of startup in the spinoff industry, which acts as a dampener on potential future cash flows. A value of $\varphi = 1$ means the spinoff is average, and $\varphi > 1$ specifies an above-average startup, with corresponding scaling to potential future cash flows.

As with our core fusion assets, we prescribe a phase progression to spinoffs. The same net present value equation holds, but we expect the spinoff phases to have a shorter threshold time to independent operation, smaller rounds of capitalization, and reduced annual cash flows. In addition, we use generalized phase progressions for each spinoff industry, rather than the asset-specific parameter values computed for fusion companies. This sufficiently captures the market characteristics absent in the determination of φ. Thus, for spinoff application s starting at state l, we have a generalized net present value equivalent to:

$$\text{NPV}_l^s = \sum_{i=l}^{\tau}\left(x_i^s * \left(\prod_{m=l}^{i} p_m^s\right) * \gamma^{i-l}\right) + \sum_{k=\tau}^{n}((X^s + \alpha^s(k-\tau)) * \gamma^k)$$

An important metric to analyze is the expected spinoff value (ESV), the discounted weighted average of cash flows from spinoffs generated by a core asset. We take the θ's to be the weighting factors, and the φ's to be the asset-specific scaling factors for the generalized net present value, giving:

$$\text{ESV}^j = \sum_{s\epsilon S}(\theta_{j,s} * \varphi_{j,s} * \text{NPV}_0^s)$$

While we can think of $\theta_{j,s}$ as the probability that company j spins off technology into market s, it is not the case that the above formula is an expectation over probability measure $\theta_{j,s}$, since $\sum_{s\epsilon S}\theta_{j,s} \neq 1$.

Based on the proposed megafund structure, the holding company will maintain the right to exclusively license the technology for a spinoff if the technological relatedness surpasses a threshold, λ, and the core fusion company does not desire to generate a divestiture on its own accord. We can evaluate the worth of this right using the restricted expected spinoff value (RESV):

$$\text{RESV}^j = \sum_{s\epsilon(S\cap\{n|\theta_{j,n}>\lambda\})}(\theta_{j,s} * \varphi_{j,s} * \text{NPV}_0^s)$$

It is trivial to note that $\text{ESV}^j \geq \text{RESV}^j$ since $s\epsilon(S \cap \{n|\theta_{j,n} > \lambda)$ is a subset of $s\epsilon S$.

For companies that have already spun-out applications, we apply a similar analysis to existing fusion companies. We take the spinoff to already exist at a certain stage, and only consider the costs, cash flows, and transition probabilities for future states. Since the megafund is unable to spin off that company's technology into the same application, we create a partition of $S = \{S^{exist}, S^{incub}$. The net present value for all spinoffs of company j is thus:

$$\sum_{s\epsilon S^{exist}}(\varphi_{j,s} * NPV_l^s) + \sum_{s\epsilon(S^{incub}\cap\{n|\theta_{j,n}>\lambda\})}(\theta_{j,s} * \varphi_{j,s} * NPV_o^s)$$

B.3 Securitization

For the general case, let there be n existing fusion companies seeking funding through a megafund that contains m debt tranches, D_i, and an equity tranche, E. We denote by $I_{tot,t}t\epsilon\{1, 2, \ldots\}$ the total investment in the megafund at time t. Similarly, let $I_{D_i,t}I_{E,t}, I_{Res,t}, t\epsilon\{1, 2, \ldots\}, i\epsilon\{1, \ldots, m\}$, represent the amount invested in the ith debt tranche, equity tranche, and reserve fund at time t, respectively. It must hold that $I_{tot,0} = I_{D_i,0} + I_{E,0} + I_{Res,0}$, but this equality no longer has to be maintained at other times, since the system is not closed (cash flows and coupon payments can cause capital to enter and leave the system). Initially, we assume that investments can only occur at the start, rather than dynamically, and thus $I_{tot,0} = I_{tot,t}$, $I_{D_i,0} = I_{D_i,t}$, and $I_{E,0} = I_{E,t}$.

The debt tranche, D_i, will have a yield of y_{D_i}, a maturity date in md_{D_i} years, and coupons will be paid out h times up to and including the maturity date. Thus, $h = 1$ corresponds to a zero-coupon bond, and $h = md_{D_i}$ corresponds to an annual coupon. Each coupon payment will therefore have a value of $\frac{md_{D_i}}{h} * y_{D_i}$. We detail return and default risk calculations in the case of a zero-coupon bond, where all debt tranches have the same maturity date. To extend this analysis to coupon bonds, we simply perform the same computation every year for each tranche's expected coupon payment after the first loss capital fund is empty.

Under these conditions, we can determine an expected year of default, which would necessitate refinancing or additional philanthropic capital.

B.4 Zero-coupon bond

A zero-coupon bond will require the payment at the maturity date of the principal of the bond and its entire yield. No payments will occur before the maturity date. We denote by $K_f (1 \leq f \leq m)$ the cash flows needed to successfully pay out the first f tranches, and compute it by $K_f = \sum_{i=1}^{f} I_{D_i,0} * (1 + y_{D_i})$. We assume that at time t, the megafund owns ρ_t^j of company j, and therefore receives $\rho_t^j * x_t^j$ of the cash flows. If company j has independently spun out its technology for spinoff application s, then the megafund will also acquire a pro rata portion of the spinoff's cash flows. Finally, we must account for the expected cash flows from the spinoff rights. If the licensing arrangement calls for company j to receive $p^{j,license}$ of all spinoffs incubated by the megafund, then the expected cash flows are the ownership-weighted RESV^j without the discount factor:

$$(1 - p^{j,license}) * \sum_{s \in (S^{incub} \cap \{n | \theta_{j,n} > \lambda\})} \left(\theta_{j,s} * \varphi_{j,s} * \sum_{i=0}^{t} \left(x_i^s * \left(\prod_{m=0}^{i} p_m^s\right)\right)\right)$$

Thus, the total expected cash flow at time t is:

$$\zeta_t^j = \rho_t^j * \left(x_t^j + \sum_{s \in S^{exist}} \left(\varphi_{j,s} * \sum_{i=l}^{t} \left(x_i^s * \left(\prod_{m=l}^{i} p_m^s\right)\right)\right)\right.$$

$$\left. + (1 - p^{j,license}) * \sum_{s \in (S^{incub} \cap \{n | \theta_{j,n} > \lambda\})} \times \left(\theta_{j,s} * \varphi_{j,s} * \sum_{i=0}^{t} \left(x_i^s * \left(\prod_{m=0}^{i} p_m^s\right)\right)\right)\right)$$

We thus expect the first f tranches, such that $f | \left(\sum_{j=1}^{n} \zeta_{md_{D_i}}^j \geq K_f, \sum_{j=1}^{n} \zeta_{md_{D_i}}^j \leq K_{f+1}\right)$, to receive their stated payouts. If all debt tranches receive their stated payouts, the remaining cash flows are allocated to the equity tranche and computing the return on investment is trivial.

Additionally, we compute the probabilities that each tranche will receive its promised return, then display how to leverage the debt coupon reserve to lower the loss probability to a level consistent with the bond ratings of credit rating agencies. In the simplest case, we assume that companies start at stage zero, and only receive cash flows following the independent operational threshold (though the analysis can easily extend to incorporate cash flows from spinoffs). In this case, the total cash flows at the maturity date are:

$$\sum_{k=\tau}^{md_{D_i}} ((X^j + \alpha^j (k - \tau)))$$

We assume that there is uncertainty for revenues after the independent operational threshold is reached, such that X^j is a random variable with mean μ^j and standard deviation σ^j. The payoff at time $t > \tau$ for company j therefore has a mean $\mu^j + \alpha^j (k - \tau)$ and standard deviation σ^j. Treating the total cash flow at the maturity date as the sum of $(md_{D_i} - \tau)$ random variables and constant terms, we derive the mean to be:

$$(md_{D_i} - \tau) * \mu^j + \alpha^j \left(\frac{(md_{D_i})((md_{D_i} + 1) - \tau(\tau - 1)}{2}\right) - \alpha^j \tau (md_{D_i} - \tau)$$

and the standard deviation:

$$\sqrt{(md_{D_i} - \tau)}\sigma^j$$

Using the same logic developed by Hull *et al.* (2019), we assume that payoffs from companies are correlated lognormal variables to avoid the negative potential payoffs arising from a multivariate normal distribution. With this total maturity date cash flow distribution, we can analyze the risk of default by computing the probability it does not exceed the promised payoff to each tranche.

Appendix C Megafund Simulation

In this section, we discuss the setup and performance of a fusion megafund portfolio using Monte Carlo simulation.

C.1 Setup of the megafund

A summary of the steps used to simulate megafund performance is included in Figure C.1. We used multi-period Monte Carlo simulation to evaluate the performance of the megafund. The boxes in blue are the parameters that serve as inputs for the simulation. These parameters comprise the probability of successful phase transition, the phase length, the cost of each phase, the probabilities of a spinoff, the correlation between separate projects, and the value of projects (i.e., the cash that can be generated by selling a project). These parameters are estimated through several different methods. Since the projects in our simulation are real fusion assets, we used data on costs accrued to date to estimate peer costs and future phase costs. Valuations and spinoff cash flows are estimated by identifying comparable assets in the energy sector. Phase transition probabilities and phase length, as well as project correlations are derived as a weighted average of several expert estimates. To account for uncertainty in these opinions, a conservative and an aggressive estimate are provided.

The first step of building a portfolio involves acquiring a fixed percentage stake in a project. Different stages of a project will require funding, and they will be financed by the megafund

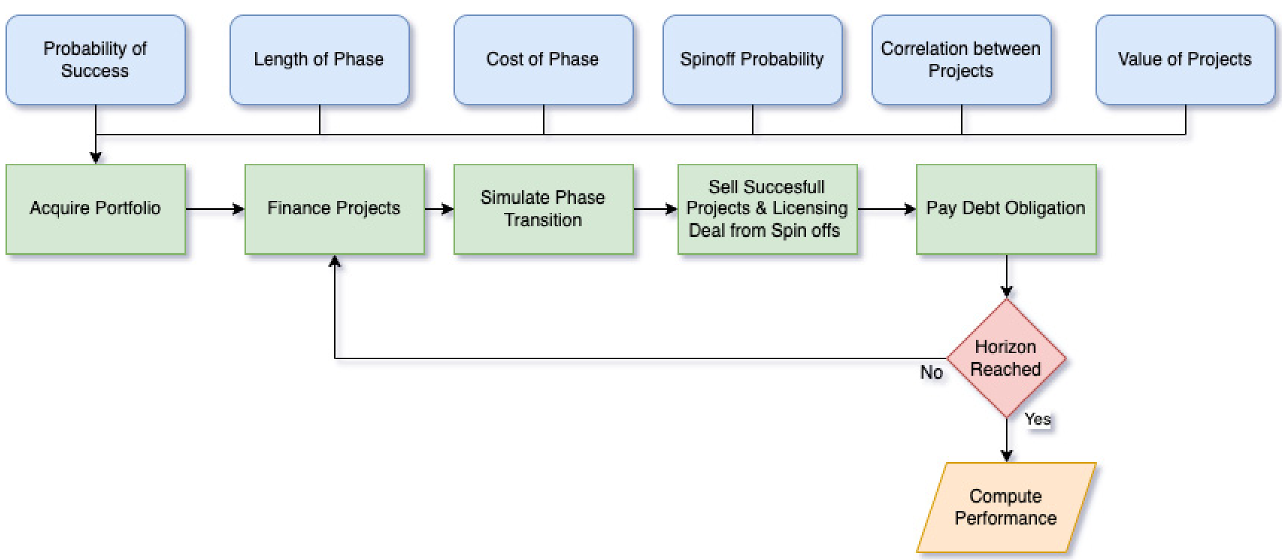

Figure C.1 Flowchart depicting the megafund simulation framework. The boxes in blue are various estimated parameters/inputs. The boxes in green depict steps of simulation.

proportional to its ownership. The sale of successful projects generates the cash flow for paying its debt obligations. Once the fund reaches its time horizon, it is liquidated, and the cash generated is used to fulfill its debt obligations and equity payments, and the performance of the fund is computed. If the fund is unable to pay the bond coupon payments at any point, it is liquidated.

C.2 Sensitivity analysis

It is important to understand the sensitivity of the megafund's performance to changes in our model's parameters. To that end, we examine the effects of varying several parameters: the cost or funding required for each stage, the duration of each stage, the probability of a successful phase transition, the revenue obtained from selling of the project at different stages and the revenue from the licensing deal for the spinoff companies.

In Table C.1, we present the performance of an equity-funded megafund after decreasing the estimated value of a successful project from \$66B to \$44B and \$22B. As expected, we find that the ROE decreases with the decreasing value of successful projects. We also observe that lowering

Table C.1 We present the equity-funded megafund portfolio performance under (a) conservative and (b) aggressive assumptions, varying the value of the successful project (revenue generated after selling a successful project) from \$66B to \$44B and \$22B. We also vary the cash flow generated from a licensing deal of a spinoff company from \$300M to \$30M.

	Revenue: \$66B	Revenue: \$44B	Revenue: \$22B	Revenue: \$22B
Measures	Spinoff: \$300M	Spinoff: \$300M	Spinoff: \$300M	Spinoff: \$30M
(a) *Conservative assumptions*				
ROE (total)	39.12	9.39	−20.34	−45.20
ROE (arithmetic avg)	1.96	0.47	−1.02	−2.26
ROE (geometric avg)	1.66	0.45	−1.13	−2.96
Pr(ROE < −10%)	3.36	42.14	78.05	88.62
Sharpe (annual)	0.11	0.04	−0.15	−0.4
# Prototyping	0.0	0.0	0.0	0.0
# Breakeven	4.2	4.2	4.2	4.4
# Commercialization	3.4	3.4	3.4	3.2
# Finished	0.1	0.1	0.1	0.1
(b) *Aggressive assumptions*				
ROE (total)	283.17	164.19	45.22	19.76
ROE (arithmetic avg)	14.16	8.21	2.26	0.99
ROE (geometric avg)	6.95	4.98	1.88	0.91
Pr(ROE < −10%)	5.45	15.44	32.27	39.77
Sharpe (annual)	0.24	0.21	0.12	0.05
# Prototyping	0.0	0.0	0.0	0.0
# Breakeven	0.1	0.1	0.1	0.3
# Commercialization	3.1	3.1	3.1	2.8
# Finished	1.2	1.2	1.2	1.2

the value of licensing deals from \$300M to \$30M while keeping the value of a successful project fixed at \$22B decreases the average annual ROE from -1.13% to -2.96%. It is apparent that the revenue from spinoff licensing deals is an important component in generating cash flows and returns for investors.

Table C.2 gives the performance of an equity-funded megafund after varying the probability of successful phase transition, the funding required at different phases, and the phase duration. As expected, increasing the probability of success leads to higher returns on equity because of fewer failures among projects. Similarly, lower costs in each phase lead to higher returns on equity. For phases of shorter duration, a greater number of projects reach their final stages during the fixed tenor of the fund, which results in higher returns on equity.

Table C.2 We present the megafund portfolio performance under (a) conservative and (b) aggressive assumptions, varying the parameters of the probability of successful phase transition (PoS), the cost of each stage, and the length of each stage of the project. The baseline equity funded megafund performance is included, and we vary the estimates from 0.8 ×baseline to 1.2×baseline, while keeping other parameters constant. The value of a successful project was fixed to \$66B and spinoff values were fixed to \$300M.

	Baseline	PoS		Cost		Length	
	1×	0.8×	1.2×	0.8×	1.2×	0.8×	1.2×
(a) *Conservative assumptions*							
ROE (total)	39.12	21.52	56.41	54.97	23.50	61.86	44.54
ROE (arithmetic avg)	1.96	1.08	2.82	2.75	1.17	3.09	2.23
ROE (geometric avg)	1.66	0.98	2.26	2.21	1.06	2.44	1.86
Pr(ROE $< -10\%$)	3.36	3.62	2.69	0.07	30.55	59.59	0.81
Sharpe (annual)	0.11	0.09	0.12	0.15	0.29	0.09	0.1
# Prototyping	0.0	0.0	0.0	0.0	0.0	0.0	2.0
# Breakeven	4.2	2.7	4.9	4.2	4.3	1.1	5.8
# Commercialization	3.4	1.7	4.9	3.4	3.2	1.4	1.7
# Finished	0.1	0.1	0.2	0.1	0.1	0.5	0.1
(b) *Aggressive assumptions*							
ROE (total)	283.2	120.9	458.6	304.5	260.7	336.2	246.0
ROE (arithmetic avg)	14.16	6.04	22.93	15.22	13.03	16.81	12.30
ROE (geometric avg)	6.95	4.04	8.98	7.24	6.62	7.64	6.40
Pr(ROE $< -10\%$)	5.45	11.48	3.98	1.90	14.05	21.60	0.10
Sharpe (annual)	0.24	0.16	0.30	0.26	0.22	0.22	0.26
# Prototyping	0.0	0.0	0.0	0.0	0.0	0.0	1.0
# Breakeven	0.1	0.0	0.2	0.0	0.3	0.0	3.6
# Commercialization	3.1	1.6	4.2	3.1	2.9	0.0	2.1
# Finished	1.2	0.5	2.0	1.2	1.2	1.8	0.9

Endnotes

1 https://www.nytimes.com/2022/12/13/science/nuclear-fusion-energy-breakthrough.html. Accessed 13 Dec. 2022.

2 https://www.iter.org. Accessed 11 Dec. 2020.

3 ITER's previous goal was 2025, but it was recently announced it will be postponed after issues at the facility emerged. See https://www.theguardian.com/science/2023/jan/06/french-nuclear-fusion-project-may-be-delayed-by-years-its-head-admits. Accessed 11 Jan. 2023.

4 Source data found at https://ourworldindata.org/energy.

5 There are other instances of major geopolitical disruptions impacting the oil economy, such as the Libyan Civil War, the Iraq War, and the Iranian revolution. Details can be found at https://abcnews.go.com/Business/supply-demand-geopolitical-tensions-oil-prices-rise/story?id=65640000. Accessed 20 Aug. 2020.

6 See https://www.worldometers.info/world-population/.

7 On February 19, 2021, the United States officially rejoined the Paris Agreement.

8 More information about the accord can be found at https://www.osti.gov/biblio/21036864-midwestern-greenhouse-gas-reduction-accord.

9 For a complete list of the largest farms, please see https://interestingengineering.com/the-11-biggest-wind-farms-and-wind-power-constructions-that-reduce-carbon-footprint. Accessed 20 Aug. 2020.

10 Additional information on the US's contribution to ITER: https://www.reuters.com/article/us-nuclear power-fusion-iter/iter-nuclear-fusion-project-avoids-delays-as-u-s-doubles-budget-idUSKBN1H2286. Accessed 05 Dec. 2020.

11 Details of the proposed American fusion reactor: https://finance.yahoo.com/news/america-finally-makes-plans-own-203200269.html. Accessed 05 Dec. 2020.

12 More details on the Mankala model and the history of the fusion industry can be accessed here: https://www.world-nuclear.org/information-library/economic-aspects/financing-nuclear-energy.aspx. Accessed 05 Dec. 2020.

13 https://www.exeltium.com/project/?lang=en#enhance-visibility-electro-intensive-industry. Accessed 05 Dec. 2020.

14 For an analysis of the economic implications of the Human Genome Project, please see: https://www.genome.gov/27544383/calculating-the-economic-impact-of-the-human-genome-project. Accessed 05 Dec. 2020.

15 Further private sector applications of the Human Genome Project: https://web.ornl.gov/sci/techresources/Human_Genome/project/privatesector.shtml. Accessed 05 Dec. 2020.

16 https://www.fusionindustryassociation.org. Accessed 10 Jan. 2023.

17 https://tae.com/fusion-power/. Accessed 8 Dec. 2020.

18 https://cfs.energy. Accessed 8 Dec. 2020.

19 https://generalfusion.com. Accessed 10 Jan. 2023.

20 https://generalfusion.com/2021/06/general-fusion-to-build-its-fusion-demonstration-plant-in-the-uk-at-the-ukaea-culham-campus/. Accessed 11 Oct. 2021.

21 https://www.tokamakenergy.co.uk. Accessed 10 Jan. 2023.

22 https://firstlightfusion.com. Accessed 8 Dec. 2020.

23 https://www.helionenergy.com. Accessed 8 Dec. 2020.

24 https://www.zapenergyinc.com. Accessed 8 Dec. 2020.

25 See report: https://www.pwc.com/gx/en/services/sustainability/assets/pwc-the-state-of-climate-tech-2020.pdf.

26 We expect that the demonstration of net energy from Lawrence Livermore National Laboratory will increase private and public funding to fusion projects.

27 See: https://www.princeton.edu/news/2017/12/15/artificial-intelligence-helps-accelerate-progress-toward-efficient-fusion-reactions. Accessed 05 Dec. 2020.

28 See: https://www2.lehigh.edu/news/nuclear-fusion-wrestling-with-burning-questions-on-the-control-of-burning-plasmas. Accessed 05 Dec. 2020.

29 See: https://3dprint.com/208394/slm-steel-nuclear-components/. Accessed 05 Dec. 2020.

30 Details of Assystem's project: https://www.neimagazine.com/news/newsassystem-wins-contracts-for-uks-step-fusion-project-7569744. Accessed 05 Dec. 2020.

31 NASA's report on the Fusion Driven rocket: https://www.nasa.gov/directorates/spacetech/niac/2012_Phase_II_fusion_driven_rocket/. Accessed 05 Dec. 2020.

32 For more details on Princeton Satellite Systems' project, please see: https://www.psatellite.com/pss-awarded-two-nasa-sbirs/. Accessed 05 Dec. 2020.

33 https://news.mit.edu/2021/MIT-CFS-major-advance-toward-fusion-energy-0908. Accessed 11 Oct. 2021.

34 For more details on PSFC's synchrocyclotron, please see https://sites.nationalacademies.org/cs/groups/bpasite/documents/webpage/bpa_062406.pdf. Accessed 10 Dec. 2020.

35 For more information on aircraft launch and recovery systems, please see: https://www.ga.com/alre. Accessed 10 Dec. 2020.

36 More fusion spinoffs, compiled by ITER, are listed here: https://www.iter.org/doc/www/content/com/Lists/Stories/Attachments/624/spinoffs.pdf. Accessed 10 Dec. 2020.

37 http://race.ukaea.uk/about-race/delivering-robotics/. Accessed 10 Dec. 2020.

38 Additional information on the 32-tesla magnet: https://nationalmaglab.org/magnet-development/magnet-science-technology/magnet-projects/32-tesla-scm. Accessed 10 Dec. 2020.

39 https://taelifesciences.com. Accessed 10 Jan. 2023.

40 Commonwealth Fusion Systems, TAE Technologies, General Fusion, Helion, Tokamak Energy, ZAP Energy, Focused Hydrogen, HB11 Energy, First Light Fusion, Type One Energy, Avalanche, Shine, LPP Fusion, MIFTi, Marvel Fusion, HelicitySpace, Hyperjet Fusion Corp, and NK Labs

41 https://www.eigpartners.com. Accessed 21 Dec. 2020.

42 The Nuclear Regulatory Commission monitors tritium-contaminated water. An overview of their measurements and mitigation techniques can be found here: https://www.nrc.gov/docs/ML0620/ML062020079.pdf. Accessed 25 Nov. 2020.

43 https://world-nuclear.org/information-library/current-and-future-generation/nuclear-fusion-power.aspx. Accessed 10 Dec. 2020.

44 https://www.newscientist.com/article/dn17952-iter-a-brief-history-of-fusion/. Accessed 25 Nov. 2020.

45 https://news.mit.edu/2016/alcator-c-mod-tokamak-nuclear-fusion-world-record-1014. Accessed 20 June 2021.

46 Release written by the Chinese Academy of Sciences: https://phys.org/news/2018-11-chinese-fusion-tool-million-degrees.html. Accessed 25 Nov. 2020.

47 https://www.iter.org/construction/construction. Accessed 10 Jan. 2023.

48 Full audit can be found at https://www.eca.europa.eu/lists/ecadocuments/jus-auditinbrief_2019/jus-auditinbrief_2019_en.pdf. Accessed 25 Nov. 2020.

49 https://www.nytimes.com/2020/09/29/climate/nuclear-fusion-reactor.html. Accessed 25 Nov. 2020.

50 ARPA-E Bethe recently announced funding for lower-cost fusion concepts such as mirrors. Their press release on the initiative can be found here: https://arpa-e.energy.gov/news-and-media/press-releases/department-energy-announces-32-million-lower-cost-fusion-concepts. Accessed 25 Nov. 2020.

51 The following *Nature* editorial outlines the status and results of Google's cold fusion research: https://www.nature.com/articles/d41586-019-01675-9.

Acknowledgments

We thank Dan Brunner, Peter Burnim, Steven Cowley, Leslie Dewan, Malcolm Handley, María Pia Iriarte, Martin Kupp, Matt Miller, MIT's Technology Licensing Office, Bob Mumgaard, Jacquelyn Pless, John Sterman, Jesse Treu, Santiago Velez, and Sam Wurzel for helpful comments and discussion, and Jayna Cummings for editorial support. We are grateful to Schmidt Futures and the MIT Laboratory for Financial Engineering for research support.

References

Barbarino, M. (2020). "A Brief History of Nuclear Fusion," *Nature Physics* **16**(9), 890–893, doi:10.1038/s41567-020-0940-7.

Bednyagin, D. and Gnansounou, E. (2011). "Real Options Valuation of Fusion Energy R&D Programme," *Energy Policy* **39**(1), 116–130, doi:10.1016/j.enpol.2010.09.019.

Brown, D. M. (2012). "Comparative Climate Change Policy and Federalism: An Overview," *Review of Policy Research* **29**(3), 322–333, doi:10.1111/j.1541-1338.2012.00562.x.

Clark, C. E. (2018). "Renewable Energy R&D Funding History: A Comparison with Funding for Nuclear Energy, Fossil Energy, Energy Efficiency, and Electric Systems R&D," *Congressional Research Service*.

Climate Action Tracker (2019). "Effect of the US Withdrawal from the Paris Agreement."

Creely, A. J. *et al.* (2020). "Overview of the SPARC tokamak," *Journal of Plasma Physics* **86**(5), 865860502, doi:10.1017/S0022377820001257.

Cui, S. *et al.* (2020). "A New Method for Improving the Tritium Breeding and Releasing Performance of China Fusion Engineering Test Reactor Phase II Helium-Cooled Ceramic Breeder Blanket," *International Journal of Energy Research* **44**(7), 5977–5989, doi:10.1002/er.5392.

DeCanio, S. J. and Fremstad, A. (2013). "Game Theory and Climate Diplomacy," *Ecological Economics* **85**, 177–187, doi:10.1016/j.ecolecon.2011.04.016.

Fagnan, D. E. *et al.* (2014). "Financing Drug Discovery for Orphan Diseases," *Drug Discovery Today* **19**(5), 533–538, doi:10.1016/j.drudis.2013.11.009.

Fernandez, J.-M., Stein, R. M., and Lo, A. W. (2012). "Commercializing Biomedical Research through Securitization Techniques," *Nature Biotechnology* **30**(10), 964–975, doi:10.1038/nbt.2374.

Fusion Industry Association (2022). *The Global Fusion Industry in 2022: Fusion Companies Survey by the Fusion Industry Association*. Washington, DC: Fusion Industry Association.

Gales, S. (1999). "Present and Future Role of Cyclotrons in Nuclear Physics and Their Applications," in *Proceedings of the 15th International Conference on Cyclotrons and their Applications, Caen, France*.

Gascon, J. C. *et al.* (2012). "Design, Challenges and Key Features for the ITER Electrical Power Distribution," *Fusion Science and Technology* **61**(1T), 47–51, doi:10.13182/fst12-a13395.

Gisler, M. *et al.* (2010). "Exuberant Innovation: The Human Genome Project," *SSRN Electronic Journal*, doi:10.2139/ssrn.1573682.

Hartwig, Z. *et al.* (2020). "VIPER: An Industrially Scalable High-Current High-Temperature Superconductor Cable," *Superconductor Science and Technology* **33**(11), 11LT01.

Herbert, G. M. J. *et al.* (2014). "A Review of Technical Issues on the Development of Wind Farms," *Renewable and Sustainable Energy Reviews* **32**, 619–641, doi:10.1016/j.rser.2014.01.055.

Hull, J., Lo, A. W., and Stein, R. M. (2019). "Funding Long Shots," *Journal of Investment Management* **17**(4), 9–41.

Intergovernmental Panel on Climate Change (IPCC) (2011). *Special Report on Renewable Energy Sources and Climate Change Mitigation*. Edenhofer, Ottmar, *et al.* (eds.), Cambridge, UK: Cambridge University Press.

International Energy Agency (IEA) (2019). "World Energy Outlook 2019," Paris. https://www.iea.org/reports/world-energyoutlook.

International Renewable Energy Agency (IRENA) (2017). "REthinking Energy 2017: Accelerating the gGLobal Energy Transformation," Abu Dhabi. https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2017/IRENA_REthinking_Energy_2017.pdf.

Jamieson, V. and Highfield, R. (2009). "ITER: A Brief History of Fusion," *New Scientist*, www.newscientist.com/article/dn17952-iter-a-brief-history-of-fusion/.

Jassby, D. (2020). "ITER Is a Showcase... for the Drawbacks of Fusion Energy," *Bulletin of the Atomic Scientists*, thebulletin.org/2018/02/iter-is-a-showcase-for-the-drawbacks-of-fusion-energy/.

Klinger, T. *et al.* (2019). "Overview of First Wendelstein 7-X High-Performance Operation," *Nuclear Fusion* **59**(11).

Kramer, D. (2018). "ITER Disputes DOE's Cost Estimate of Fusion Project," *Physics Today*, physicstoday.scitation.org/do/10.1063/PT.6.2.20180416a/full/.

Lamb, W. F. and Minx, J. C. (2020). "The Political Economy of National Climate Policy: Architectures of Constraint and a Typology of Countries," *Energy Research & Social Science* **64**, 101429, doi:10.1016/j.erss.2020.101429.

Liberman, M. A. *et al.* (1999). *Physics of High-Density Z-Pinch Plasmas*. New York, NY: Springer Science, library/current-and-future-generation/nuclear-fusion-power.aspx.

Lo, A. W. and Siah, K. W. (2021). "Financing Correlated Drug Development Projects," *Journal of Structured Finance* **27**(1), 17–34.

Masuda, T. *et al.* (2007). "Development of 3-in-One High-Temperature Superconducting Cable," *Electrical Engineering in Japan* **162**(1), 15–24, doi:10.1002/eej.20577.

Miller, L. M. and Keith, D. W. (2018). "Observation-Based Solar and Wind Power Capacity Factors and Power Densities," *Environmental Research Letters* **13**(10), 104008, doi:10.1088/1748-9326/aae102.

Mukoyama, S. *et al.* (2018). "Superconducting Joint of REBCO Wires for MRI Magnet," *Journal of Physics: Conference Series* **1054**.

National Academies of Sciences, Engineering, and Medicine (NASEM) (2021). *Bringing Fusion to the US Grid*. Washington, DC: The National Academies Press. https://doi.org/10.17226/25991.

Philibert, Cédric (2006). *Barriers to Technology Diffusion: The Case of Solar Thermal Technologies*, IEA, Paris. https://www.iea.org/reports/barriers-to-technology-diffusion-the-case-of-solar-thermal-technologies. Accessed 21 Nov 2021.

Phillips, J. E. and Easterly, C. E. (1980). "Sources of Tritium," *Oak Ridge National Laboratory Union Carbide*.

Ritter, S. K. (2016). "Cold Fusion Died 25 Years Ago, but the Research Lives on," *C&EN*, cen.acs.org/articles/94/i44/Cold-fusion-died-25-years.html.

Romanov, A. *et al.* (2020). "High Frequency Response of Thick REBCO Coated Conductors in the Framework of the FCC Study," *Scientific Reports* **10**(1), doi:10.1038/s41598-020-69004-z.

Sepulveda, N. A. *et al.* (2018). "The Role of Firm Low-Carbon Electricity Resources in Deep Decarbonization of Power Generation," *Joule* **2**(11), 2403–2420, doi:10.1016/j.joule.2018.08.006.

Sun, X. (2019). "Solar Technology Got Cheaper and Better in the 2010s. Now What?" *Greentech Media*, Greentech Media, www.greentechmedia.com/articles/read/solar-pv-has-become-cheaper-and-better-in-the-2010s-now-what.

Teller, E. (ed.). (1981). *Fusion, Vol. 1: Magnetic Confinement (Part A)*. New York, NY: Academic Press.

Timilsina, G. *et al.* (2014). "A Review of Solar Energy Markets, Economics and Policies," *Solar Energy*, 167–215, doi:10.1201/b17731-12.

Tollefson, J. (2021). "US Achieves Laser-Fusion Record: What It Means for Nuclear-Weapons Research." *Nature* **597**, 163–164, https://doi.org/10.1038/d41586-021-02338-4.

Weinberg, A. M. (1994). *The First Nuclear Era: The Life and Times of a Technological Fixer*. New York, NY: AIP Press.

Wu, P.-I. *et al.* (2014). "Climate Game Analyses for CO_2 Emission Trading among Various World Organizations," *Economic Modelling* **36**, 441–446, doi:10.1016/j.econmod.2013.08.025.

Xu, Y. (2016). "A General Comparison between Tokamak and Stellarator Plasmas," *Matter and Radiation at Extremes* **1**(4), 192–200, doi:10.1016/j.mre.2016.07.001.

Keywords: Fusion investments; green investments; fusion megafund; fusion portfolio; de-risking strategy; decarbonization; financial engineering.