DCVC

001

Deep Tech Opportunities Report

Augmenting Human Abilities
Climate Tech
Computational Biology and Health

2023

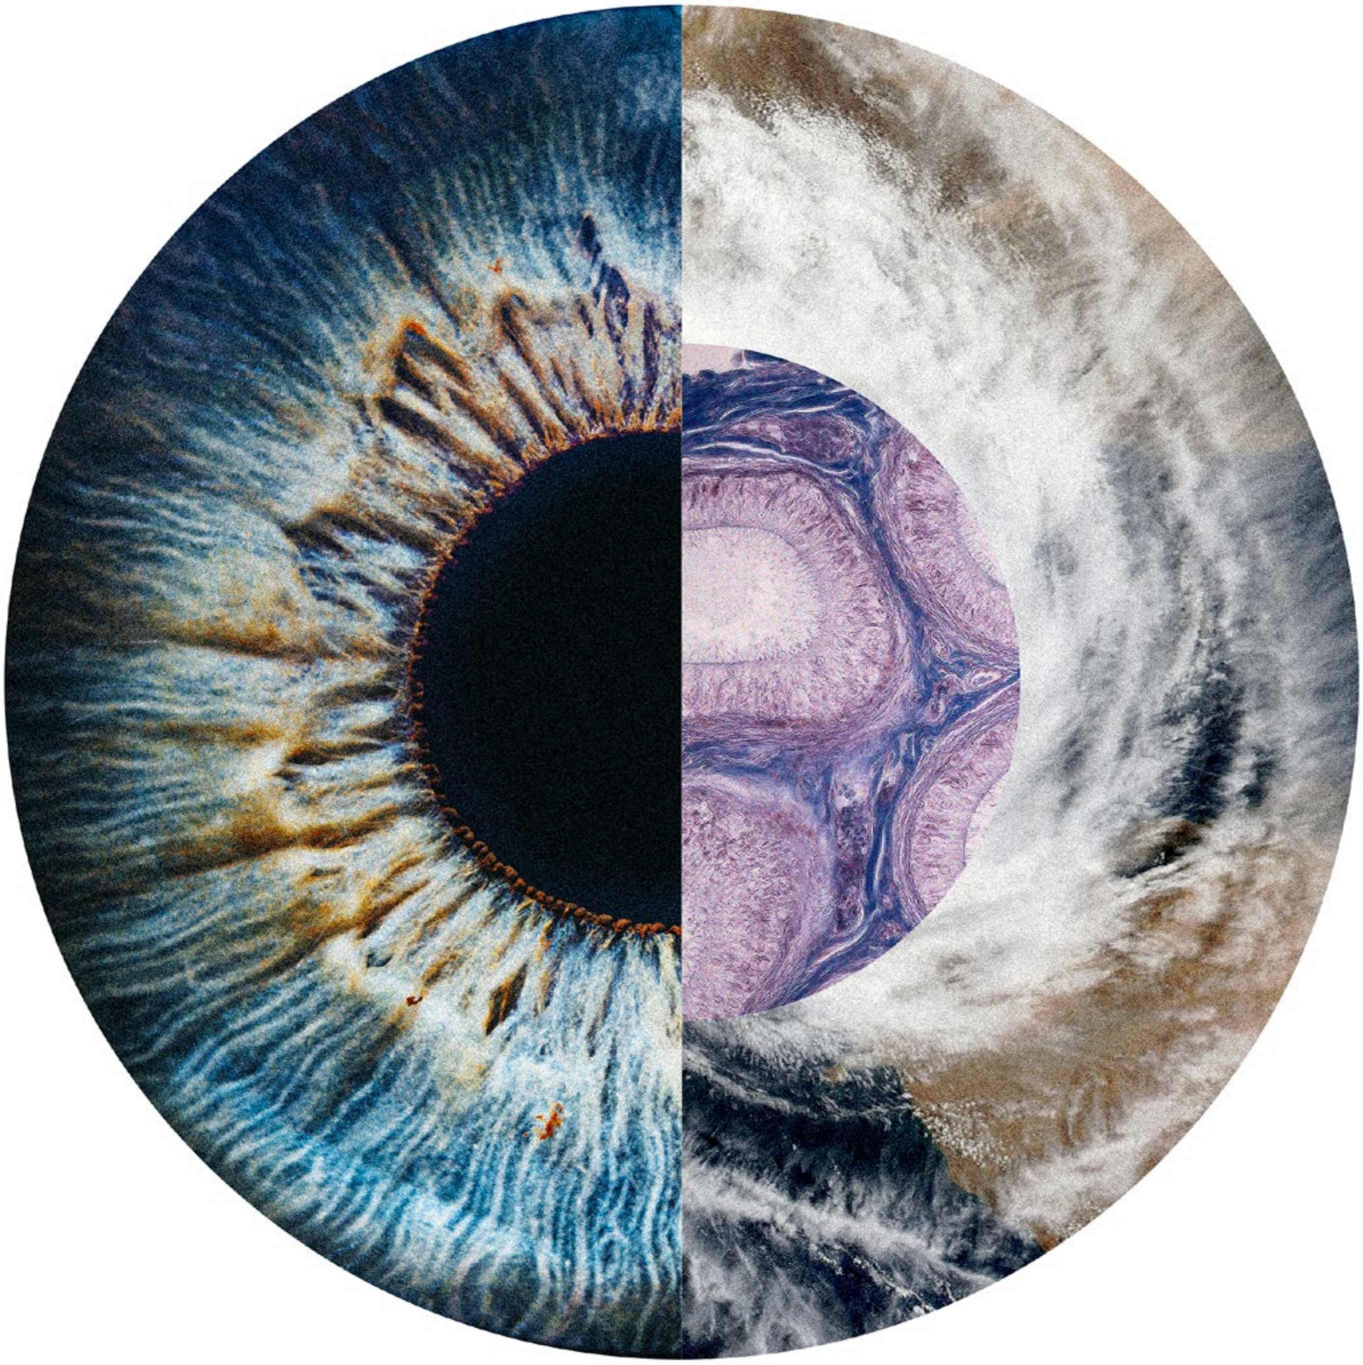


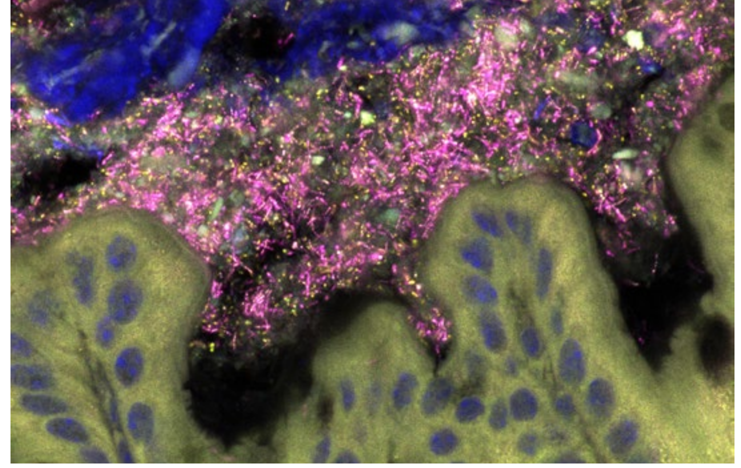



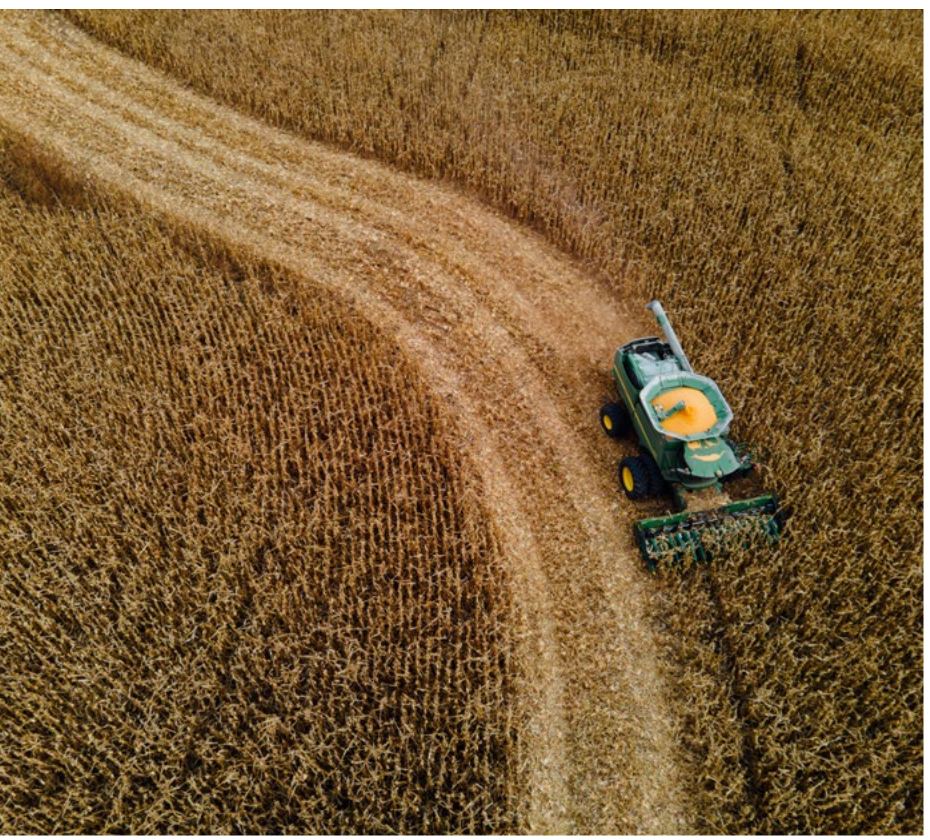


lumafield

DCVC is deep tech venture capital.

We back entrepreneurs who use emerging computational techniques to advance breakthroughs in science and technology, solving hard, real-world, trillion-dollar problems in industries that have resisted change. These are problems that matter to everyone, and that when addressed assiduously by founders of uncommon talent become immense opportunities.

Energy, water, chemicals, agriculture, mining, medicines, healthcare, space—to varying degrees, and with exceptions, these fields have been marked for decades by mostly incremental innovations. As different as they are, they share one trait: they have been at least as concerned with atoms as they have been with bits.

There's a perception that radical reinvention of inherently physical industries can come only at outrageous capital cost, and over long, unknowable timescales. But it *is* possible to disrupt these industries in ways that drastically reduce both capital expenditures and operating expenses, and to do it all at a cost and over a time period appropriate for venture firms like ours.

The secret, we believe, is to employ machine vision, robotics, simulation, large high-value datasets, and machine learning (along with other forms of AI) to attack old science and engineering problems from new angles. Only in the last decade has computing matured to the point that it can accelerate new ideas in physical and life sciences into broad commercial use. Today, it is speeding up the design-build-test-learn cycle and reducing resource requirements in every field from energy to materials science to drug development.

Look at Zap Energy, an energy startup in Washington state. Founded just six years ago, the company used simulations to show that a phenomenon called sheared flow stabilization Z-pinch could, in theory, be used to confine the plasma in a fusion reactor without the gigantic and expensive magnets other fusion companies are building. Last summer, around the same time we invested, the company converted theory to reality, feeding 500 kiloamps (kA) of current through a plasma for the first time—just 150 kA shy of the level needed to achieve Q=1, also known as scientific energy breakeven. "To be a practical energy source, we need to go well beyond Q=1," says Zap's president, Benj Conway. "But if you want to get fusion on the grid in time to make a difference to the planet, then the ability to iterate quickly on a small, cheap platform is absolutely vital."

The truth is that when the growing power of computation is applied to the world of atoms, roadblocks that felt immovable suddenly begin to crumble. We can start seeing ways to program machines that free humans from the most tedious or backbreaking forms of work; to reconfigure manufacturing, construction, and distribution chains around clean electricity; to build a new generation of safe, carbon-free generation facilities to provide that electricity; to treat and recycle enough fresh water for a growing world population; to increase crop yields without polluting our land and water; and to quickly develop cures for rare, previously untreatable genetic disorders and the hundreds of separate diseases we call cancer.

That's the difference deep tech can make and must make. Some suggest we should simply give up and adapt to trends like widening economic inequality and a catastrophically warming climate. But innovation boosted by computation is the key to mastering those trends and expanding human possibilities.

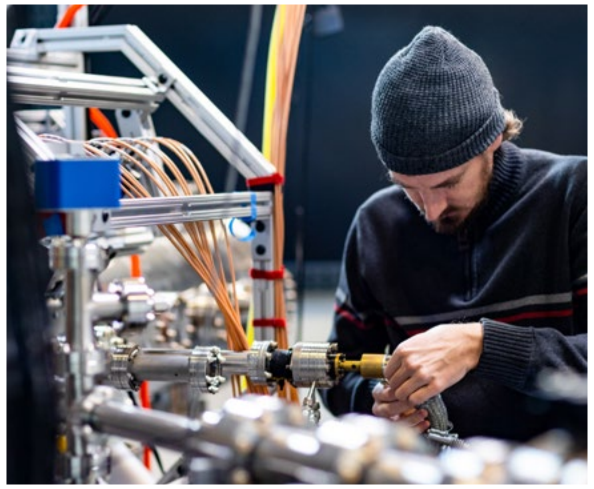

▸ **Zap Energy** is building a prototype fusion reactor that confines a filament of heat-generating plasma within a tight tube-shaped magnetic field.

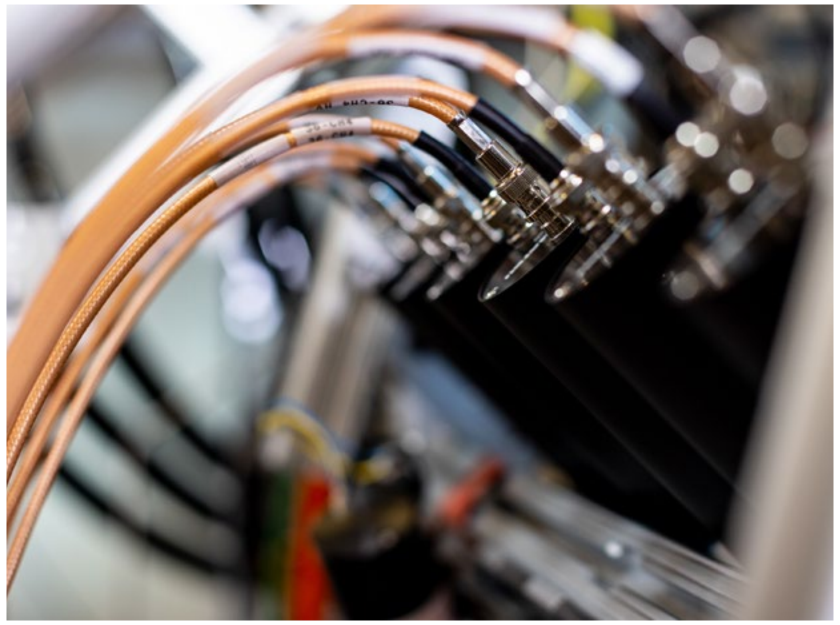

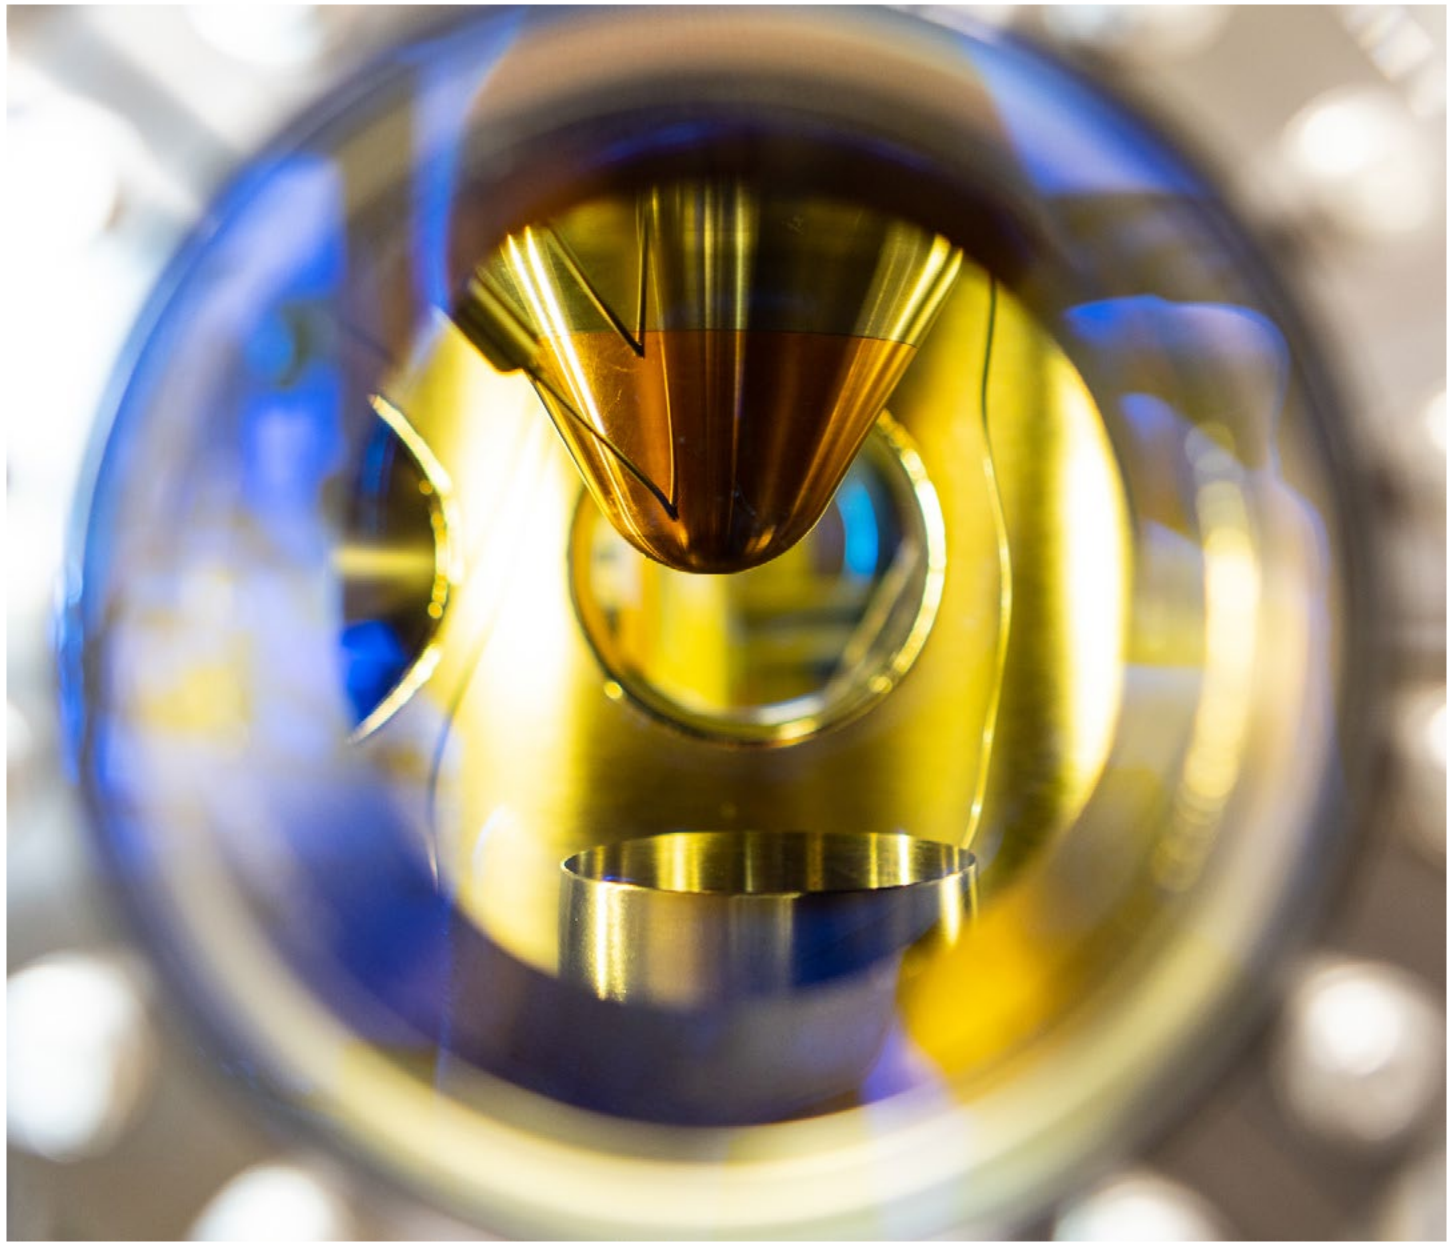

Zap stabilizes the plasma using a phenomenon called sheared flow stabilization, which looks a little like highway traffic moving at different speeds in different lanes. Zap began by modeling the technique in computer simulations.

On the DCVC investing team, bench skills, industry experience, and published papers count for more than MBAs. That means our partners have an exceptional sense of which technology opportunities will be most transformative, which are closest to market, and which entrepreneurs are best qualified to get them there. This Deep Tech Opportunities Report showcases our partners' insights and opinions about nine specific prospects, across three broad themes: augmenting human abilities (Page 12), tackling climate change (Page 34), and transforming healthcare and agriculture (Page 62). We also touch on our work in aerospace, algorithmic finance, and an area we call carbon-intelligent materials.

This isn't an exhaustive survey of DCVC's portfolio. We invest in many companies outside these areas, and in future reports we will have more to say. But the contentions we share here illustrate our thinking about the deep tech investment areas we consider the most urgent and most promising in 2023—along with a few we consider premature or outright misguided.

The optimism that's a prerequisite for deep tech investing doesn't lead us to downplay the difficulty of the problems our entrepreneurs are trying to solve. It simply has us biased toward action—toward getting to work on the technologies that will lengthen healthy lifespans, redirect centuries of fossil-fueled growth toward something sustainable, and provide food, water, jobs, and education for a near-future population of 10 billion.

"The way is certainly both short and steep," W. H. Auden once wrote, "however gradual it looks from here." It's time to start up the slope. The companies we describe here are doing that, and it's our privilege to help them. Opportunities, as described in what follows, abound.

"The way is certainly both short and steep," W. H. Auden once wrote, "however gradual it looks from here." It's time to start up the slope. The companies we describe here are doing that, and it's our privilege to help them. Opportunities, as described in what follows, abound.

Zachary Bogue & Matt Ocko

Founders and Managing Partners, DCVC

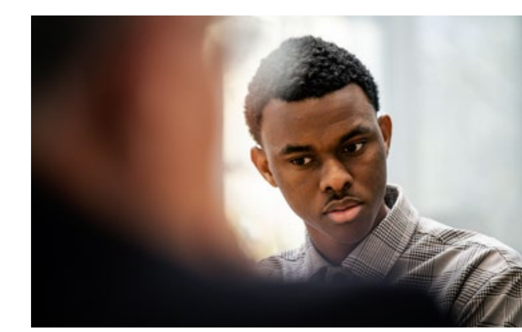

About DCVC

Founded in 2010, DCVC invests in deep tech companies from early-stage startups to established growth businesses. Our investments span agricultural technology, applied AI, climate technology, computational biology and healthcare, finance and Web3, foundational technology, industrial transformation, security, and space. DCVC currently manages more than $3 billion in committed capital across 12 funds. DCVC backs entrepreneurs applying deep tech tools—AI, big data, computation, cloud infrastructure, and applications of molecular and biological science—to transform giant industries and solve trillion-dollar problems.

2023

Deep Tech Opportunities Report

This report summarizes DCVC's thinking about the deep tech investment areas we consider the most exciting, important, and consequential in 2023. It's also a guide to the inspiring work innovators inside and outside the firm's portfolio are doing to extend human capabilities, save the environment, and make everyone's lives longer, healthier, and easier.

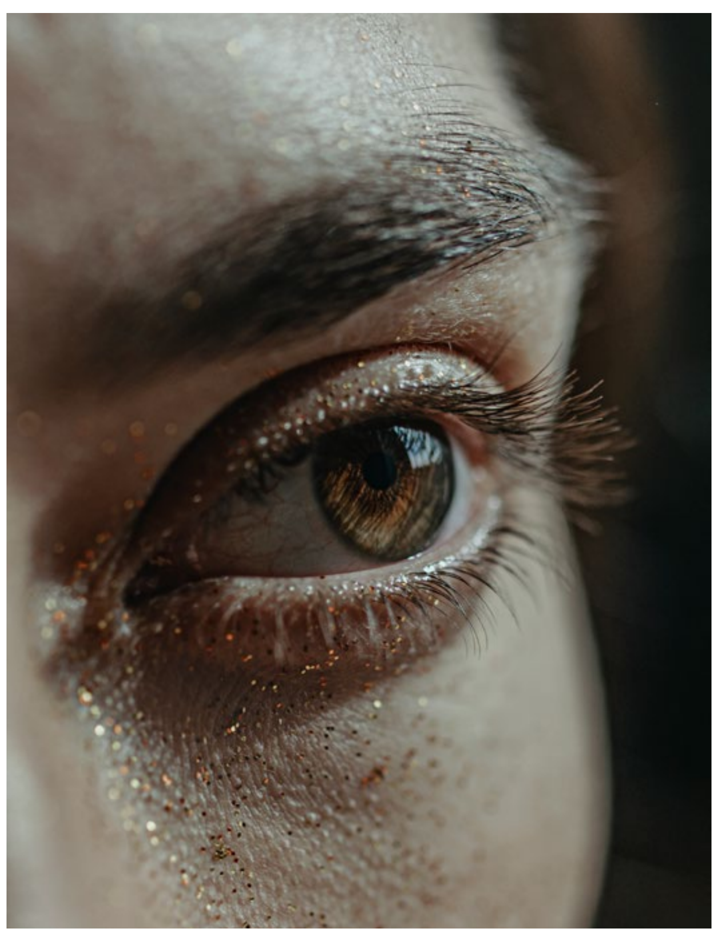

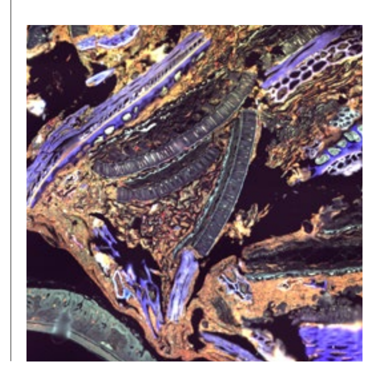

1.0

Augmenting Human Abilities

We view computing as the core technology enabling new breakthroughs across all deep tech, from manufacturing to materials to carbon-free energy to personalized medicine. But the impact of computing is currently most visible at the human-machine boundary, where robots and algorithms are taking over ever more of the tasks that have hitherto burdened people.

Humans should be free to do what humans do best: express themselves creatively, devise strategies and plans, teach and care for one another, organize with others to achieve larger goals. Wherever possible, therefore, we should ask machines to handle the most repetitive, demanding, or enfeebling forms of labor.

But it's only recently, thanks to breakthroughs in areas such as robot design and machine learning, that innovators have found ways to extend automation beyond the factory floor. Robots are now showing up in less regimented workplaces and on our streets and highways. Machine-vision technology is allowing computers to go beyond simple number-crunching and engage in much more sophisticated forms of perception and analysis. Neural-network architectures such as large language models, acting on Internet-scale data, are learning how to mimic human art and human language.

Computer scientists aren't on the cusp of creating artificial general intelligence; that's still years or decades away, if it will ever be possible. But they are teaching machines how to do more of the jobs that people can't, wont, or simply shouldn't do. Here are three contentions that explain DCVC's deep commitment to R&D in these fields, along with a comment about one hard problem where we don't see the necessary innovations on the horizon.

► Robots will finally work alongside humans, in human spaces.

Featured Companies

Agility Robotics [DCVC]

Augmenting Human Abilities: 3D Space

In factories and warehouses, robots have long been caged in safe zones where they can't harm humans. But advances in robotics software and design are gradually allowing more types of robots to operate safely *alongside* humans, in settings built *for* humans. That will help companies realize cost savings and accelerate the transfer of dull, dirty, and dangerous jobs from human workers to machines.

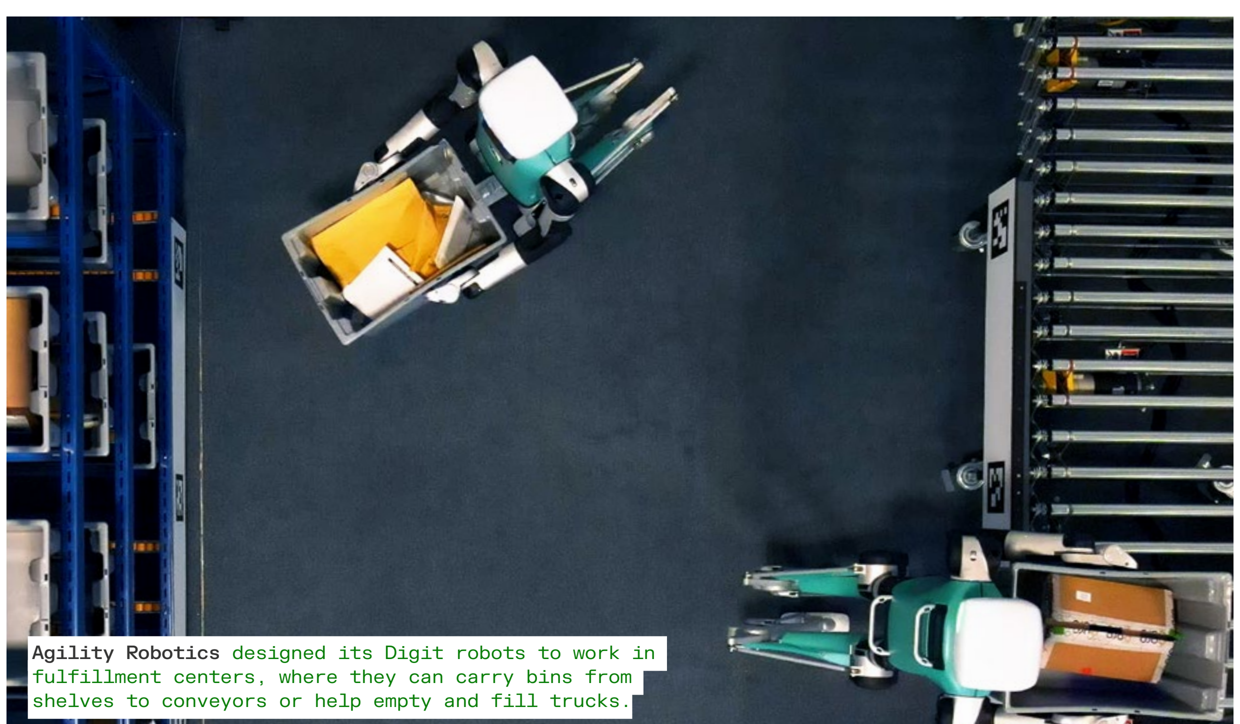

Agility Robotics designed its Digit robots to work in fulfillment centers, where they can carry bins from shelves to conveyors or help empty and fill trucks.

Today's industrial robots, such as those on most automobile assembly lines, do their jobs with implacable determination and speed. The barrier guards and sensor-enabled safety zones in these factories are there to keep human workers out of the machines' way. But an alternative is to program robots to respond to their environments, rather than smash through them along rigid, predetermined paths. That's the guiding belief behind the work of Agility Robotics, a maker of humanoid robots that's been part of the DCVC portfolio since 2020.

Cofounder and CTO Jonathan Hurst spent the first decade of his career studying the dynamics of locomotion in animals and robots. That led him to deemphasize traditional *position control*—under which a robot arm or end effector is programmed to move to a precise set of *x*, *y*, and *z* coordinates regardless of what might be in the way—in favor of *force control*, where sensor feedback helps guide the arm while keeping its torque within safe limits.

Robots built to work alongside people should be just as "floppy" as humans are, Hurst says. "We don't get to impose upon the world the action we want, the way an industrial robot arm does," he notes. "Compliant interaction and force-sensitive behavior is really important for all forms of physical interaction. We build everything else on top of those dynamical foundations."

> "
> ## We need to fill the unfillable jobs with robots. There aren't enough people.
>
> **Zachary Bogue**
> Managing Partner, DCVC

Agility's flagship robot, Digit, is 5 feet 2 inches tall and shaped like a human, except for its stubby head and long, birdlike feet. It's optimized to pick up and carry boxes and bins in a setting such as a warehouse, a fulfillment center, or the back of a semitrailer. Because its limbs operate under force control, it walks with a confident, animallike gait, adapting instantly to bumps, curbs, ramps, or rough terrain that would fell traditional robots.

Digit is a multipurpose, human-centric robot built to automate existing warehouses without costly retooling.

Digit weighs 100 pounds, so you might not want it to tip over on top of you, but otherwise it's unlikely to injure a human being. "We've never had a person hurt by the robot," says Agility's CEO, Damion Shelton. In fact, when Digit sees a person in its path, it will politely step aside.

Agility is on a mission to replace human labor in demonstrably short supply, without having to rebuild existing work environments to separate robots and humans. In the warehousing sector, the demand for labor vastly exceeds the supply of willing workers, thanks in part to high turnover. In January 2023 alone, 218,000 people in the U.S. quit their jobs in the transportation, warehousing, and utilities sector, according to the Bureau of Labor Statistics, leaving 628,000 jobs open. "The main challenge from the employer's perspective is that people don't like doing the work," says Shelton. He hopes companies will give their most physically demanding jobs to Digits, while steering humans toward other jobs they enjoy and find meaningful.

"We need to fill the unfillable jobs with robots," says DCVC managing partner Zachary Bogue. "There aren't enough people." And if those robots are built to operate in the spaces where humans already work, then—to borrow an Agility company slogan—they'll ultimately enable humans to be more human.

▸ Computers will increasingly see and react to visual patterns that humans can't perceive.

Featured Companies

Smartex [DCVC]
Plotlogic [DCVC]
Recycleye [DCVC]
Blue River Technology [DCVC]
Verdant Robotics [DCVC]
Recursion [DCVC]
Strateos [DCVC]

Augmenting Human Abilities: Machine Vision

Digital cameras can gather visual data at higher speed and resolution than the human eye, and machine-learning algorithms can spot patterns in that data with higher accuracy and consistency. That's allowing game-changing advances in every industry where imagery can inform decision-making.

Computers don't perceive the visual world the way humans do. In many ways, they have a decided advantage.

- A human eye constantly flits, or saccades, back and forth to place the fovea—the high-resolution part of the retina—on objects of interest, but a computer sees an entire visual field evenly, out to the edges of its sensors.
- The human retina is limited to the visual spectrum, but digital detectors can see far into infrared and ultraviolet wavelengths and beyond, or use lasers and time-of-flight calculations to construct a precise 3D model of a space.
- The human brain's visual centers often infer a story to make up for noisy, incomplete, or changing visual data, but a computer sees what's actually there, all the time, and is faster to recognize and sort shapes and patterns.
- A human observer will eventually get tired or distracted, but a computer can fix anything with a relentless stare.
- The human brain's visual memory is sketchy, impressionistic, and prone to fabrication, but a computer can store each frame of sensor data precisely and forever.

This last difference may have the most impact, because it's now possible to train startlingly powerful machine-learning algorithms on raw pixel data—which can be obtained at low cost and enormous scale.

These are some of the factors at work in machine vision, a technology revolution powering startups in areas as diverse as mining, agriculture, textiles, healthcare, and drug development. While the Agility robots mentioned above do employ depth cameras and light detection and ranging (lidar), that company is focused on solving the problems that arise when machines try to move through 3D spaces used by humans. Here, we're talking about a different group of companies that are combining sensor data and machine learning to *find patterns in the world that humans could never spot.*

Consider textile making, one of humanity's oldest trades. Today, most knitted fabrics are created on room-sized circular knitting machines that draw in yarns from hundreds of separate spools and knit them together automatically. The process is fast, but defects are common, and human inspectors don't always catch them. Smartex builds arrays of high-resolution cameras that record every inch of a fabric roll as it emerges from a knitting machine. Machine-learning software scans the images for recurring defects such as lines or stripes, allowing operators to stop the machine before the whole production run is marred.

Smartex says its ultimate goal is not just to automate textile factories, but to help eliminate the 92 million tons of textile waste they generate each year. "With computer vision, we're talking about continuous monitoring and a kind of accuracy that we aren't able to see with the naked eye," says DCVC managing partner Matt Ocko, who has decades of experience working with companies using AI on hard, physical problems. He predicts that Smartex's technology will help ease labor shortages, not just replacing the labor of one worker with one machine, "but supercharging it, 1-to-100 or 1-to-1,000."

Or consider another ancient industry, mining. Each rock type absorbs or reflects a signature set of wavelengths of visible, infrared, or ultraviolet light, so when mining companies want to know which walls of a pit or which piles of rubble contain valuable ore, they can send samples to the lab for spectral analysis. But the turnaround time can be hours or days. A DCVC-backed Australian startup called Plotlogic is automating the process and moving it on site. The company's portable cameras fuse 3D measurements (Continued on page 24)

Verdant Robotics
Agriculture

Verdant uses computer vision to target fertilizer application and laser weeding—making farming greener while saving chemicals and labor.

Strateos
Drug discovery

Biopharma companies using Strateos' Cloud Lab platform can conduct experiments remotely, speeding up the design, synthesis, testing, and analysis of new molecules.

Recycleye
Waste management

Recycleye's robots scan a moving waste stream and use machine learning to sort out high-value recyclable plastics from material that will be downcycled or landfilled.

Recursion
Drug discovery

Recursion treats arrays of cells with fluorescent dyes and analyzes them under high-throughput microscopes to see how drugs, toxins, pathogens, or genetic alterations change their growth.

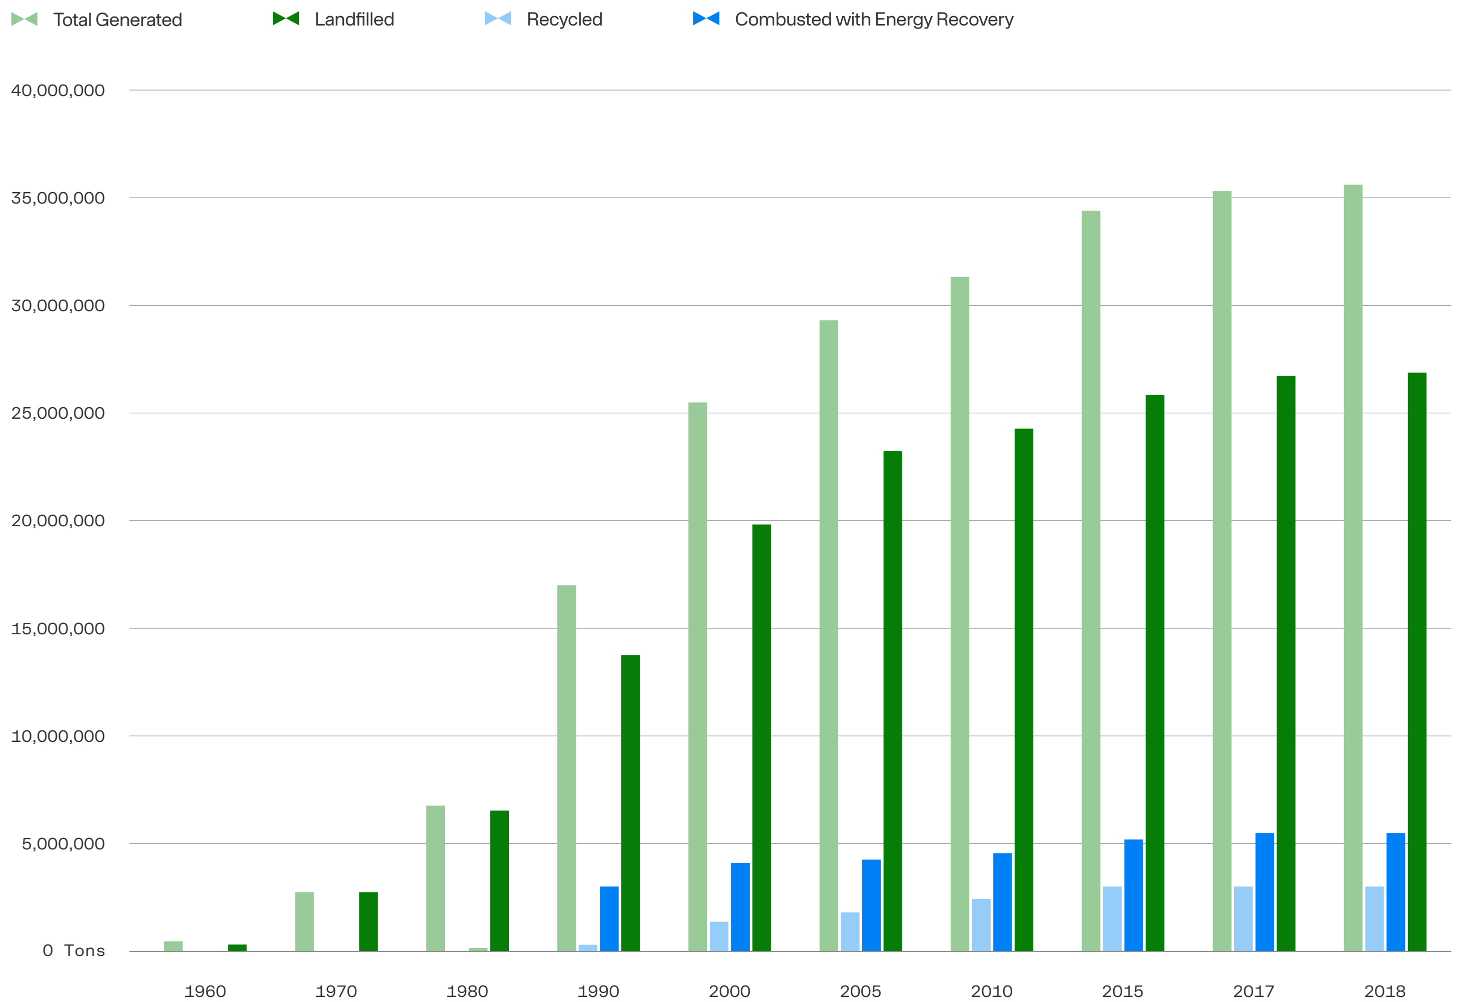

Figure 1.2 U.S. Plastics Waste "Management"

The vast majority of plastic waste generated in the U.S. gets landfilled rather than recycled or burned for energy. Computer-vision-equipped robots could lower the cost of sorting mixed-waste streams to help divert more plastic for recycling.

Source: EPA

from lidar sensors with hyperspectral feedback, and use machine-learning algorithms to assess rocks in real time—fast enough to guide a power shovel or spot poor-quality ore as it passes on a conveyor belt. The same idea works in recycling plants, where sensor-equipped robots like those built by DCVC-backed startup Recycleye can scan a waste stream on a conveyor belt and sort out the high-value recyclable plastics from material that will be burned, downcycled, or landfilled.

On farms, meanwhile, computer-vision algorithms can be trained to tell crops from weeds. Blue River Technology, a DCVC-backed company that was bought by John Deere in 2017, built such a system into tractors, and today Deere's "See & Spray" system helps farmers apply herbicides only where they're needed. Another startup backed by DCVC Bio, Verdant Robotics, is taking that idea to the next level, using computer vision to target fertilizer delivery and guide a laser weeding system.

Drug-discovery companies and health-tech providers are putting software to work to spot subtle variations in digital images. Recursion, in Salt Lake City, trains high-throughput microscopes on arrays of cells treated with fluorescent dyes to see how different toxins, pathogens, genetic changes, and small and large molecules change the cells' morphology over time. This phenomics data can show whether a candidate drug is having the intended effect at the cellular level. Strateos, another DCVC-backed company, also uses computer vision in its remotely operated "cloud labs" to speed up the design, synthesis, testing, and analysis of new drug molecules.

Outside of DCVC's portfolio, other companies are using machine vision to improve on human perception in areas like radiology, pathology, surveillance, geoanalytics, and disaster response. "If anything, I think computer vision is underhyped as a breakthrough," says James Hardiman, a DCVC general partner who has led many of the firm's investments in this area. "Anything where a person is looking, and trying to render a decision or judgment, is ripe now for augmentation."

Putting computer vision to work in patient care

Physicians, nurses, and medical technicians rely heavily on visual information. But no health professional can spot everything. Several DCVC companies are giving them the tools to see more.

Caption Health

Removing a key barrier to cardiac preventive care by allowing medical technicians to perform echocardiograms

SafelyYou

Analyzing real-time video to detect falls in dementia-care facilities

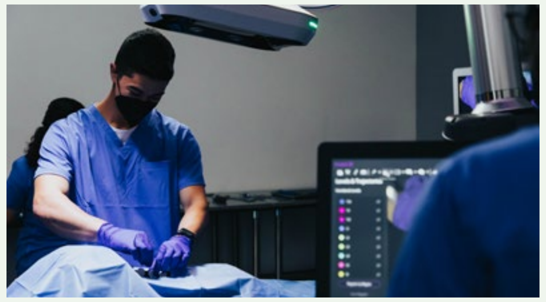

Proprio

Using light-field technology to help surgeons navigate during spinal surgery

Swift Medical

Assisting in wound care with smartphone-based digital planimetry

▸ Generative AI will disrupt every business activity built around communication.

Featured Companies

OpenAI
Google
Microsoft

Augmenting Human Abilities: Large Language Models

As computers learn to synthesize content—using machine-learning systems to generate sophisticated texts, computer code, images, and video— every business activity built around search, discovery, and communication will be disrupted.

We've long had to constrain the way we talk to computers to the narrow formats and topics they're programmed to understand. The rule-based expert systems of the 1970s and 1980s, which deduced information from a fixed base of knowledge, were brittle and uncomprehending outside their specific domains. 1990s-era search engines fixated on keywords (and returned only the most popular or most promoted results, not the most appropriate ones). But in 2022, software designed to answer open-ended queries in natural language finally crossed a critical threshold of power and usability, allowing us to communicate with computing systems as if they were almost human. The replies generated by the newest large language models, such as OpenAI's GPT-4 and Google's LaMDA, are so literate, detailed, and plausible that it's possible—for the first time—to imagine how generative AI might take over or substantively augment some forms of work.

"Almost everything companies do is underpinned by written communication," says Hardiman, who trained in engineering physics and sits on the boards of numerous early-stage AI startups. "The ability of machines to generate human-level sentences and paragraphs—that's a really powerful and disruptive technology." Add large language models' ability to generate executable software code, and even to solve design problems in fields such as chemistry, biology, and materials science, and you get a potential automation revolution in the routine aspects of research and knowledge work, perhaps allowing humans to focus their creativity elsewhere.

The turning point arrived in the form of machine-learning models called transformers, running on purpose-built supercomputers powerful enough to hold hundreds of gigabytes of text snippets or parameters in active memory. OpenAI's language model GPT-4 (for Generative Pre-trained Transformer 4) can respond to natural-language prompts by guessing what a human would write. It's able to summarize texts, translate between languages, beat most humans at standardized exams, give step-by-step reasoning for its answers to complex questions, and take on creative challenges such as crafting lyrics, poetry, jokes, and movie pitches. OpenAI's models aren't doing what we would all call reasoning or thinking—they work by statistically parsing thoughts already recorded by humans, in the form of hundreds of billions of words from sources such as Wikipedia, published books, and an archive of web data. But they do produce answers with the polish and sheen of intelligence.

Never before has software been capable of producing convincing, extended responses to unconstrained prompts. The quality of musings from GPT-4's predecessor, GPT-3, startled and even alarmed many observers when OpenAI let developers and reporters try out the model beginning in 2020—New York Times columnist Farhad Manjoo called it "at once amazing, spooky, humbling and more than a little terrifying." The larger Internet community expressed similar awe when OpenAI released a user-friendly, dialogue-optimized version of GPT-3 as ChatGPT in November 2022. Because the model had ingested so much data, it could expound on almost anything. It could counsel a videogame addict on their career options, offer a detailed postmortem analysis of the Challenger disaster, or rewrite a recipe for pineapple upside-down cake in the style of Shakespeare. GPT-4 is even more capable: it outscores most humans on exams such as the LSAT and the GRE, and can turn a crude napkin-style drawing of a website into working HTML and Javascript.

These drastic leaps in the capabilities of large language models were unexpected, and they seem to have been a product of the models' increased parameter capacity, as well as the amount of data used to train them. (GPT-3's pre-training dataset was 10 times as large as that used by the second-biggest model at the time—and yet it was quickly surpassed by models from companies such as Nvidia, Microsoft, and Google's DeepMind, and its own successor, GPT-4.)

The larger these models grow, the more skilled they become, emergently acquiring the ability to do math, answer physics questions, or string together arguments in ways that seem to mimic common-sense reasoning. It's one of the most vivid examples to date of what Google AI researchers once called the "unreasonable effectiveness of data": the advantage that statistical models have over elaborate hand-coded rules when it comes to emulating human thought. "So far, these models just get better and more general with more parameters," says DCVC operating partner Steve Crossan, who previously led AI product development work at DeepMind and GlaxoSmithKline. There is no known theoretical limit to how far this trend can go, Crossan says—though, in practice, there may be a ceiling on the cost of processing and the availability of training data. But every area where DCVC invests is already being affected, Crossan says, and from the firm's current perspective—we write this in May 2023—the impact can only grow.

Today's large language models still have limits. Before they can provide useful answers, they require extensive fine-tuning by humans, in part through a process called reinforcement learning from human feedback. Users can prise out better answers if they know a little about how the models work—a burgeoning field called prompt engineering. And because large language models look for statistical patterns in language—not for truth or consistent meaning—they're prone to confabulation, producing what their creators call hallucinations and what outside observers have called "fluent bullshit." But none of these limitations seem to be inherent to the technology. And we expect many of them to fade as time goes on and as AI companies invest more in work on interpretability (making it easier to understand how the models arrive at their decisions) and alignment (ensuring the systems follow the intentions of their designers).

It's already clear that large language models will have a range of workplace applications, starting with contexts where humans will have the final say. The models can generate first drafts of marketing copy, website text, and blog and social media posts. They can invent names for businesses or products. They will likely take over much content moderation and paralegal work. And they will make it possible to automate entire genres of routine business writing where form is more important than flair—think contract drafting, vendor negotiations, online tutoring, and chat- or email-based customer service. OpenAI says it uses GPT-4 internally for support, sales, content moderation, and programming, "with great impact" in each area.

"The first problems this is going to be well-suited for are low-stakes tasks that still require human oversight, but where the models speed up the execution," says Hardiman. "But a lot of companies are also starting to iterate and experiment on where else you can take this technology, and what other kinds of problems you can apply it against."

One of those is writing software. OpenAI's Codex tool uses a descendant of GPT-3 that has also been trained on billions of lines of Python code (which is, after all, just another form of language). Microsoft worked with OpenAI to incorporate Codex into its GitHub software development platform, where it can rough out sections of code in response to comments written by human programmers. Rather than writing code themselves (or copying it from Stack Overflow), a developer can simply describe what they want, and Codex will generate a first draft of the working code. As capabilities like this spread to more software development platforms, "most programming will become code reviewing," Crossan predicts.

Strikingly, transformer networks trained on natural language and small amounts of scientific data can produce meaningful answers to questions about organic chemistry, gene expression, and protein and RNA structure, beating state-of-the-art machine-learning approaches developed specifically for these applications. These models may be gleaning from their data some generalized model of the world; more likely, they're just stochastic parrots, echoing patterns they've detected without connecting them to meaning. Either way, they promise to help generate novel chemistries and biomolecular designs, and to speed up many forms of research. We also expect impressive and convincing results from multimodal AIs—those that both accept and emit text, images, and video. (GPT-4 is already multimodal, though some of those features aren't currently turned on in ChatGPT.) "That's going to have a big impact on the collective psyche when it hits," says Crossan.

The creators of large language models have a responsibility to build in safeguards to keep them from being used for mischief. Humans will need to edit, rework, fact-check, and stress-test their output to filter out the biases and misinformation they're likely to pick up from their training data. Alignment researchers will need to find ways to keep bad actors from using AI to generate deepfakes and other disinformation for large-scale influence operations. And employers, educators, and politicians will need to think ahead about how to value and reward human labor and creativity in a future where large language models—and other generative AI tools in areas like art and music—have automated some forms of work.

All that said, many rote and boilerplate forms of text require minimal creativity; handing them over to software might spare human brainpower for the tasks that truly reward it. A future with deeper and more meaningful forms of human-computer collaboration would benefit everyone. We see the recent work of OpenAI and its partners and competitors as a potential step in that direction.

> "The ability of machines to generate human-level sentences and paragraphs—that's a really powerful and disruptive technology.
>
> James Hardiman
> General Partner, DCVC

Figure 1.3
Sizes of large language models

The sophistication of a large language model depends mainly on the number of parameters built into the model—that is, the number of values that can be adjusted during training to fine-tune its performance. The first version of the ChatGPT chatbot used a version of OpenAI's GPT-3, which had more parameters than any other large language model at the time of its release in 2020, but was soon surpassed by models from Google and Microsoft. GPT-4 is larger than GPT-3, but OpenAI has not revealed how much larger.

Source: Hisamoto, S. (2022, June 12). *Sizes of large language models*. Observable. Retrieved from observablehq.com/@sorami/sizes-of-large-language-models.

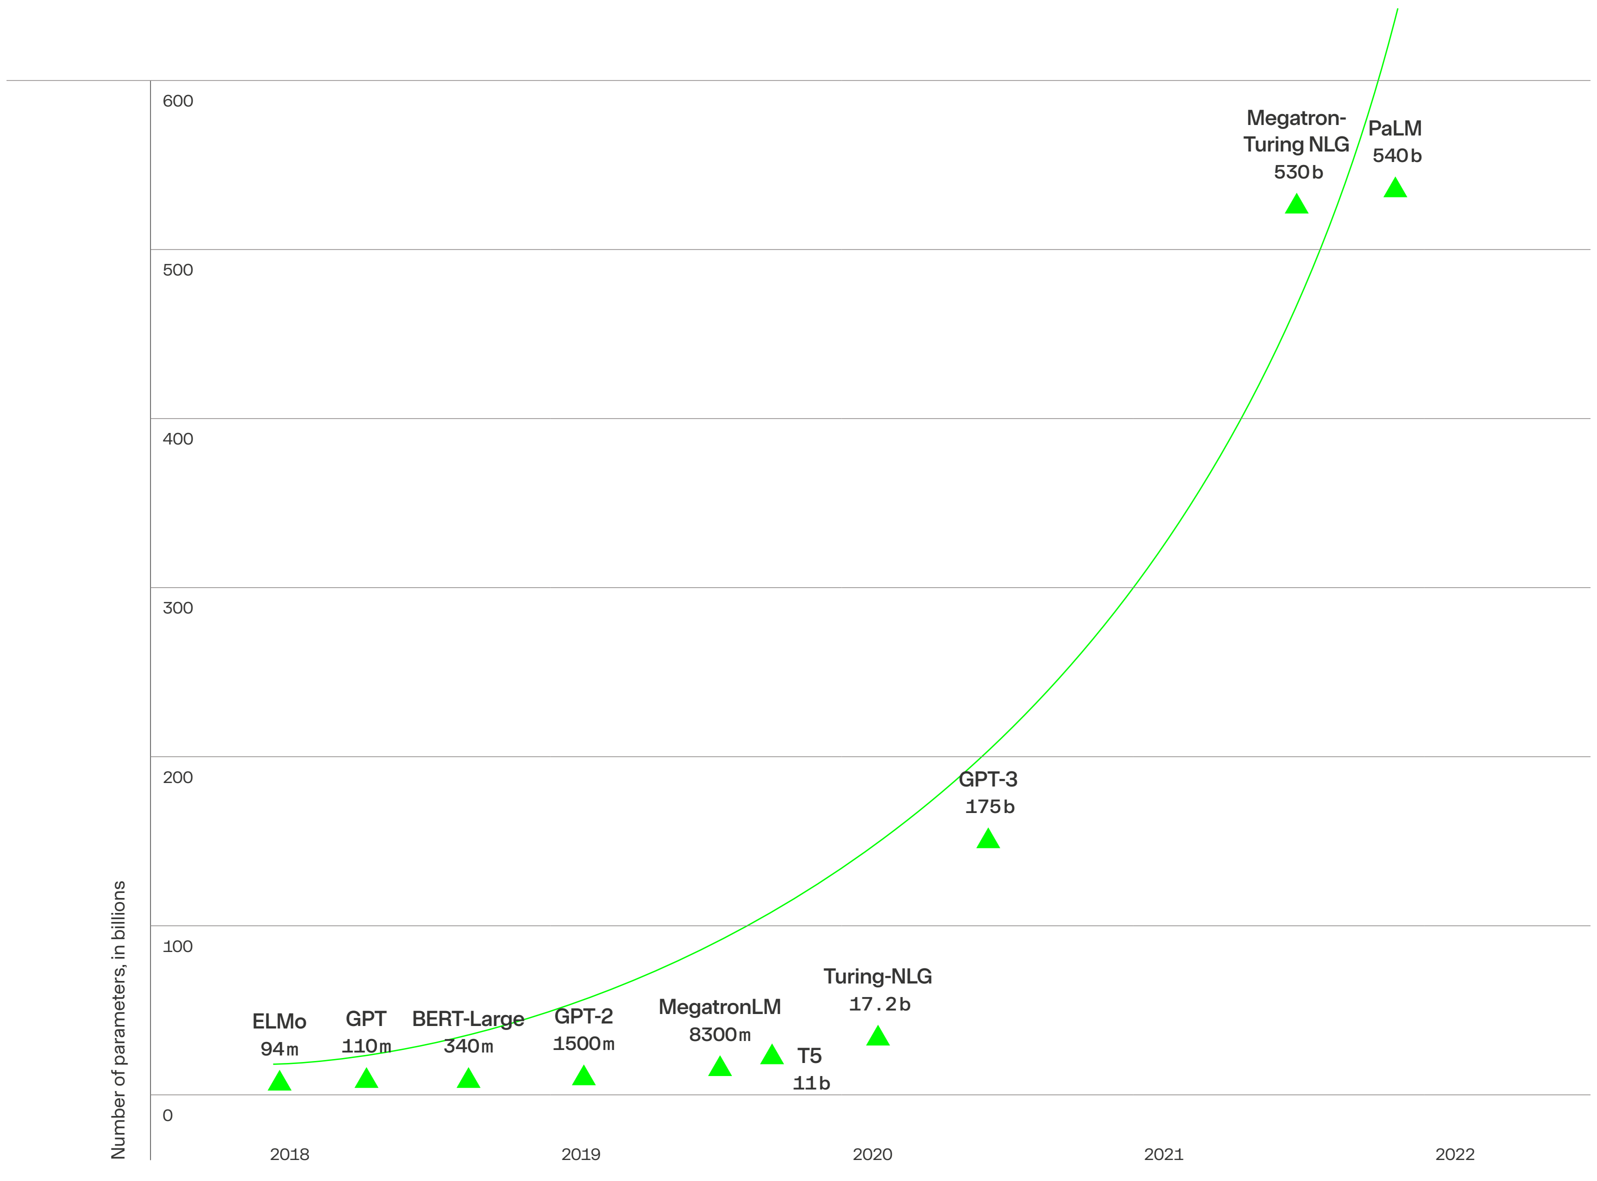

Speech Interfaces

Though they can understand a few formulaic commands, voice-driven virtual assistants remain as uncomprehending as infants. The problem is that spoken conversation, unlike text-based chat, is a wickedly unstructured domain.

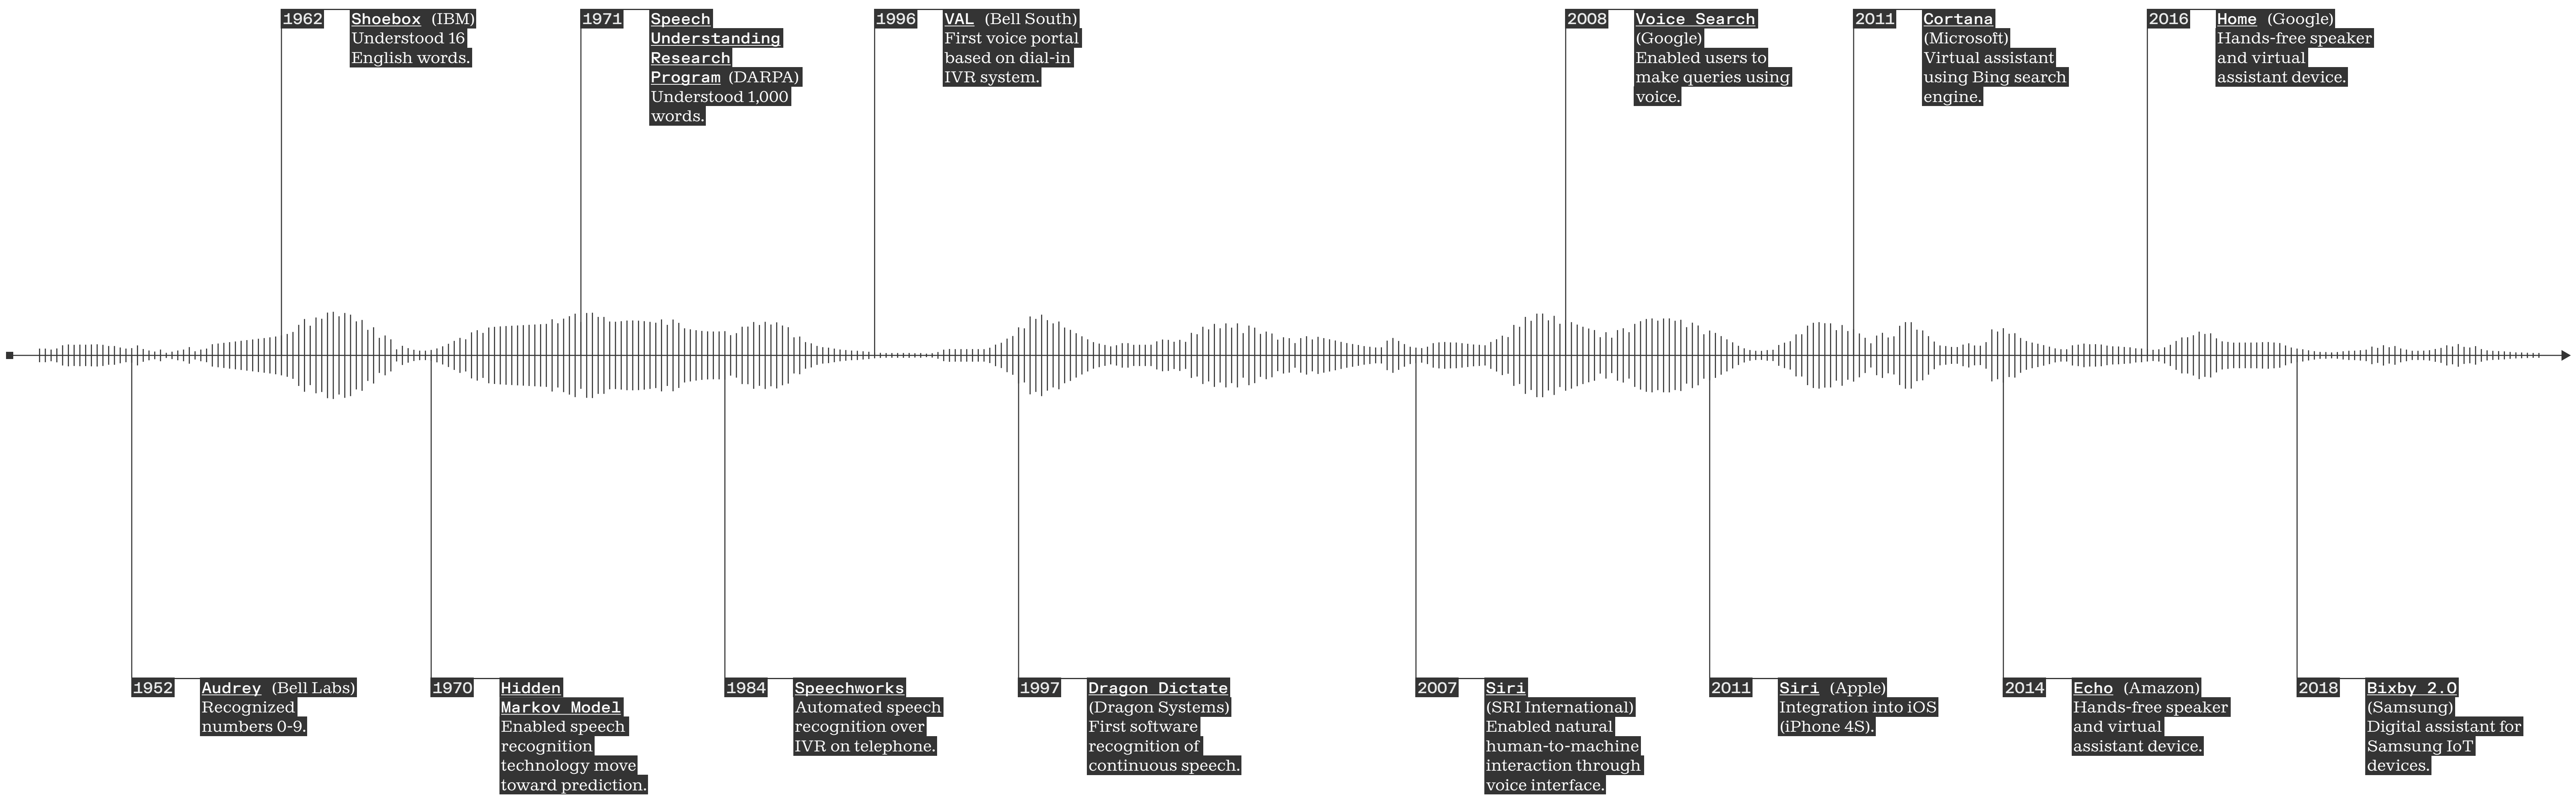

Figure 1.4 | The Evolution of Voice-driven Interfaces

Source: Apexon

↳ This report is focused on the biggest present opportunities in deep tech investing. But we will occasionally talk about a few technologies we might categorize as "enticing, but unready to advance at the pace required by venture capital." These aren't dumb ideas—just hard problems where we don't see the required technical insights or market demand on the near horizon, or where fundamentally new thinking is required.

The first of these is the challenge of interacting with our computers through speech, as science fiction characters have been doing for decades. Existing interactive voice response (IVR) technology works well enough in brief, highly structured situations—say, arranging a mail-order prescription drug refill by phone. But the more variables an interaction involves, the faster these systems frustrate human users.

The central problem is that most transactions in our daily lives involve too many options to account for in a scripted, programmed way. In automated customer-support calls, for instance, we quickly zero out to an operator because we know that only a human will be able to respond flexibly or ask the questions that get to the heart of the matter. "A computer wants to provide a framework—a structure like dates and times—and it's painful to try to describe the constraints and then ask you to specify those constraints by voice," says Hardiman. "I'm skeptical that voice interfaces are superior to text outside of a handful of applications."

It's nice to be able to tell a virtual assistant like Siri or Alexa to give you a weather forecast, play a song, relay a sports score, turn on your lights, schedule an appointment, or add an item to a shopping list. But what's notable is how little the core capabilities of these programs have evolved over the years. Siri, Alexa, Cortana, and Google Home are far from being able to track the context that makes a real conversation coherent, or handle the routine yet improvisatory tasks we might trust to a human personal assistant, like sending out party invitations, picking a vacation spot, or researching the options in a medical crisis.

"The excitement around chatbots illustrates how Siri, Alexa and other voice assistants...have squandered their lead in the AI race," the New York Times observed in March 2023. Voice assistants are still "dumb as a rock," Microsoft CEO Satya Nadella admitted to the Financial Times. (Amazon, for one, seems to be only reducing its investment in this area; a wave of layoffs in 2022 reportedly focused on the Alexa division.)

Companies will surely try building some of the features of chatbots into a Siri- or Alexa-style assistant. (It's already possible to implement a custom shortcut that connects Siri to ChatGPT, allowing users to query OpenAI's system via voice and see the answer displayed on the screen.) But even if these technologies fully converge, we don't think this is where large language models will shine. The predictive models that power chatbots can sift through millions of web pages and decades of transcripts to synopsize what humans have said or written, but they still can't genuinely *converse*, because conversation is mysteriously, irreducibly social and unstructured. Also, because these models are grounded in probabilities, not in common-sense reasoning or real-world factual knowledge, they still sometimes cheerfully and unknowingly tell you things that are just not true—which is seldom helpful in a conversation partner or a helper.

But the more fundamental problem is that speech is a such a unidimensional, low-bandwidth way to convey information. It could turn out be the least imaginative way to communicate with our AI helpers—even if that flies the face of decades of science fiction representations of talking computers and robots. "Voice is just a crappy interface for doing lots of tasks, like booking plane tickets," says Hardiman. "I don't want a machine to read out to me all the options of various flights—it's information that should clearly be laid out and interacted with visually."

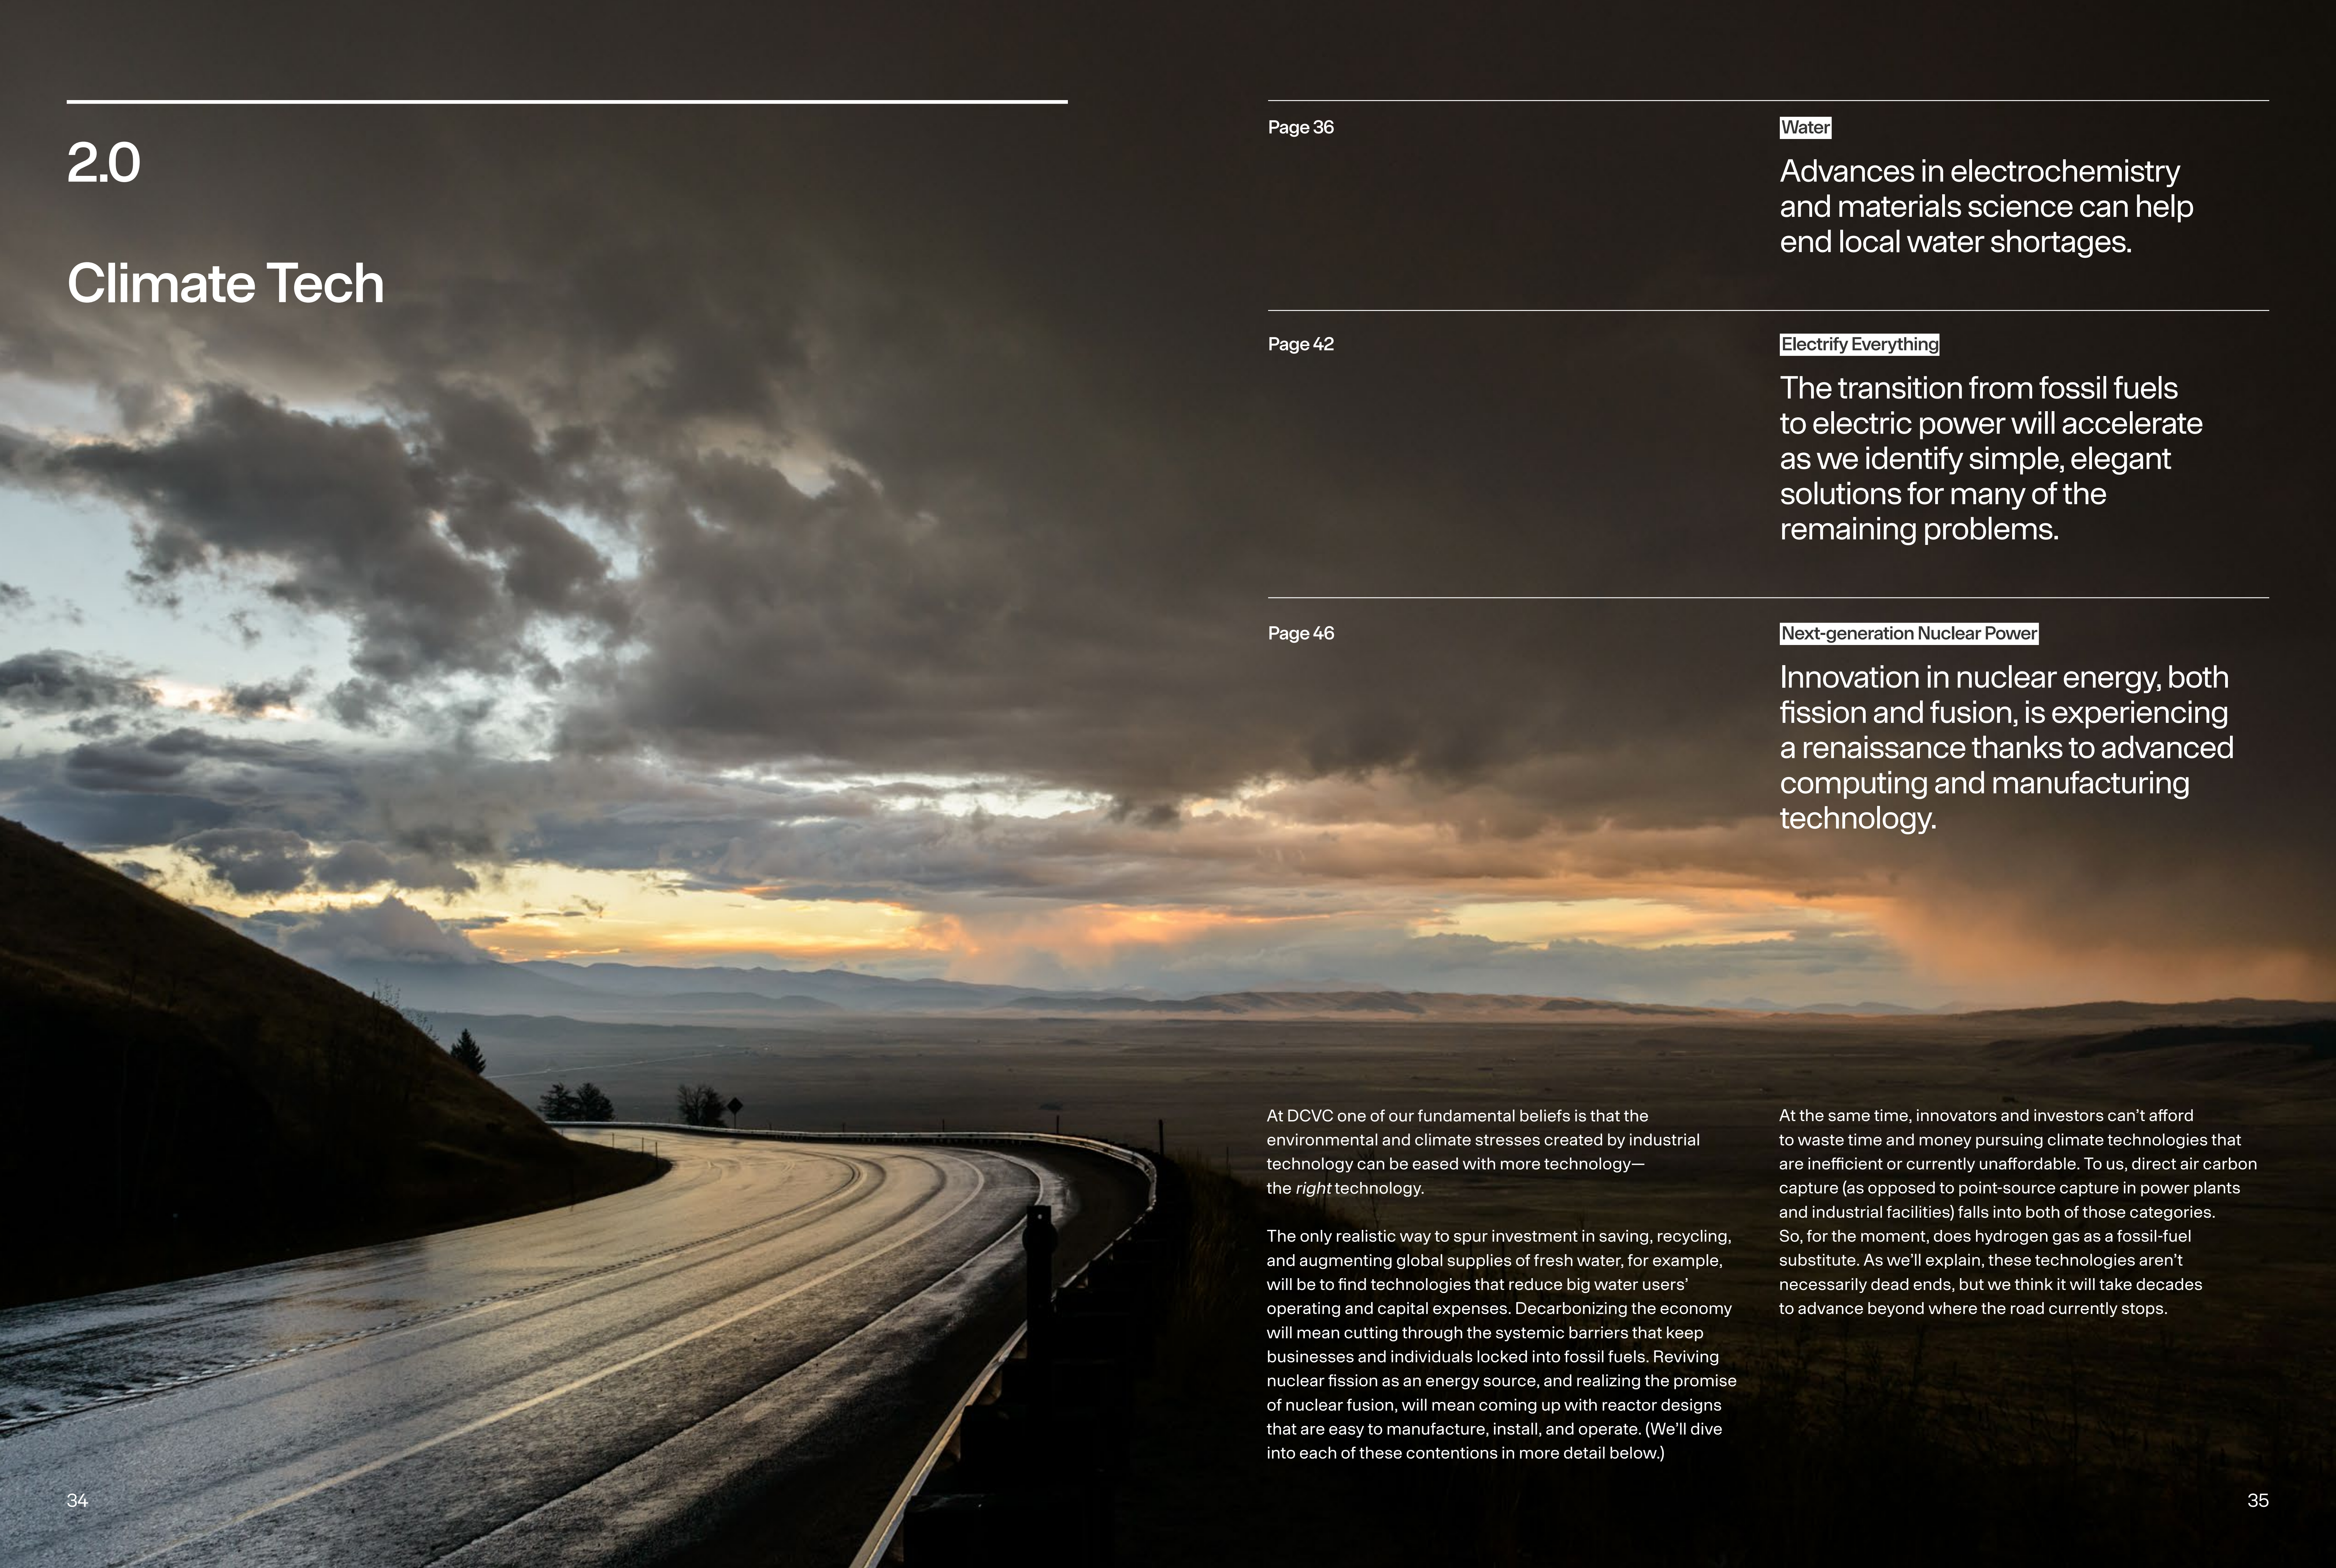

2.0

Climate Tech

At DCVC one of our fundamental beliefs is that the environmental and climate stresses created by industrial technology can be eased with more technology—the *right* technology.

The only realistic way to spur investment in saving, recycling, and augmenting global supplies of fresh water, for example, will be to find technologies that reduce big water users' operating and capital expenses. Decarbonizing the economy will mean cutting through the systemic barriers that keep businesses and individuals locked into fossil fuels. Reviving nuclear fission as an energy source, and realizing the promise of nuclear fusion, will mean coming up with reactor designs that are easy to manufacture, install, and operate. (We'll dive into each of these contentions in more detail below.)

At the same time, innovators and investors can't afford to waste time and money pursuing climate technologies that are inefficient or currently unaffordable. To us, direct air carbon capture (as opposed to point-source capture in power plants and industrial facilities) falls into both of those categories. So, for the moment, does hydrogen gas as a fossil-fuel substitute. As we'll explain, these technologies aren't necessarily dead ends, but we think it will take decades to advance beyond where the road currently stops.

▶ Advances in electrochemistry and materials science can help end local water shortages.

Featured Companies

Aclarity [pcvc]
Aquafortus [pcvc]
ZwitterCo [pcvc]

Climate Tech: Water

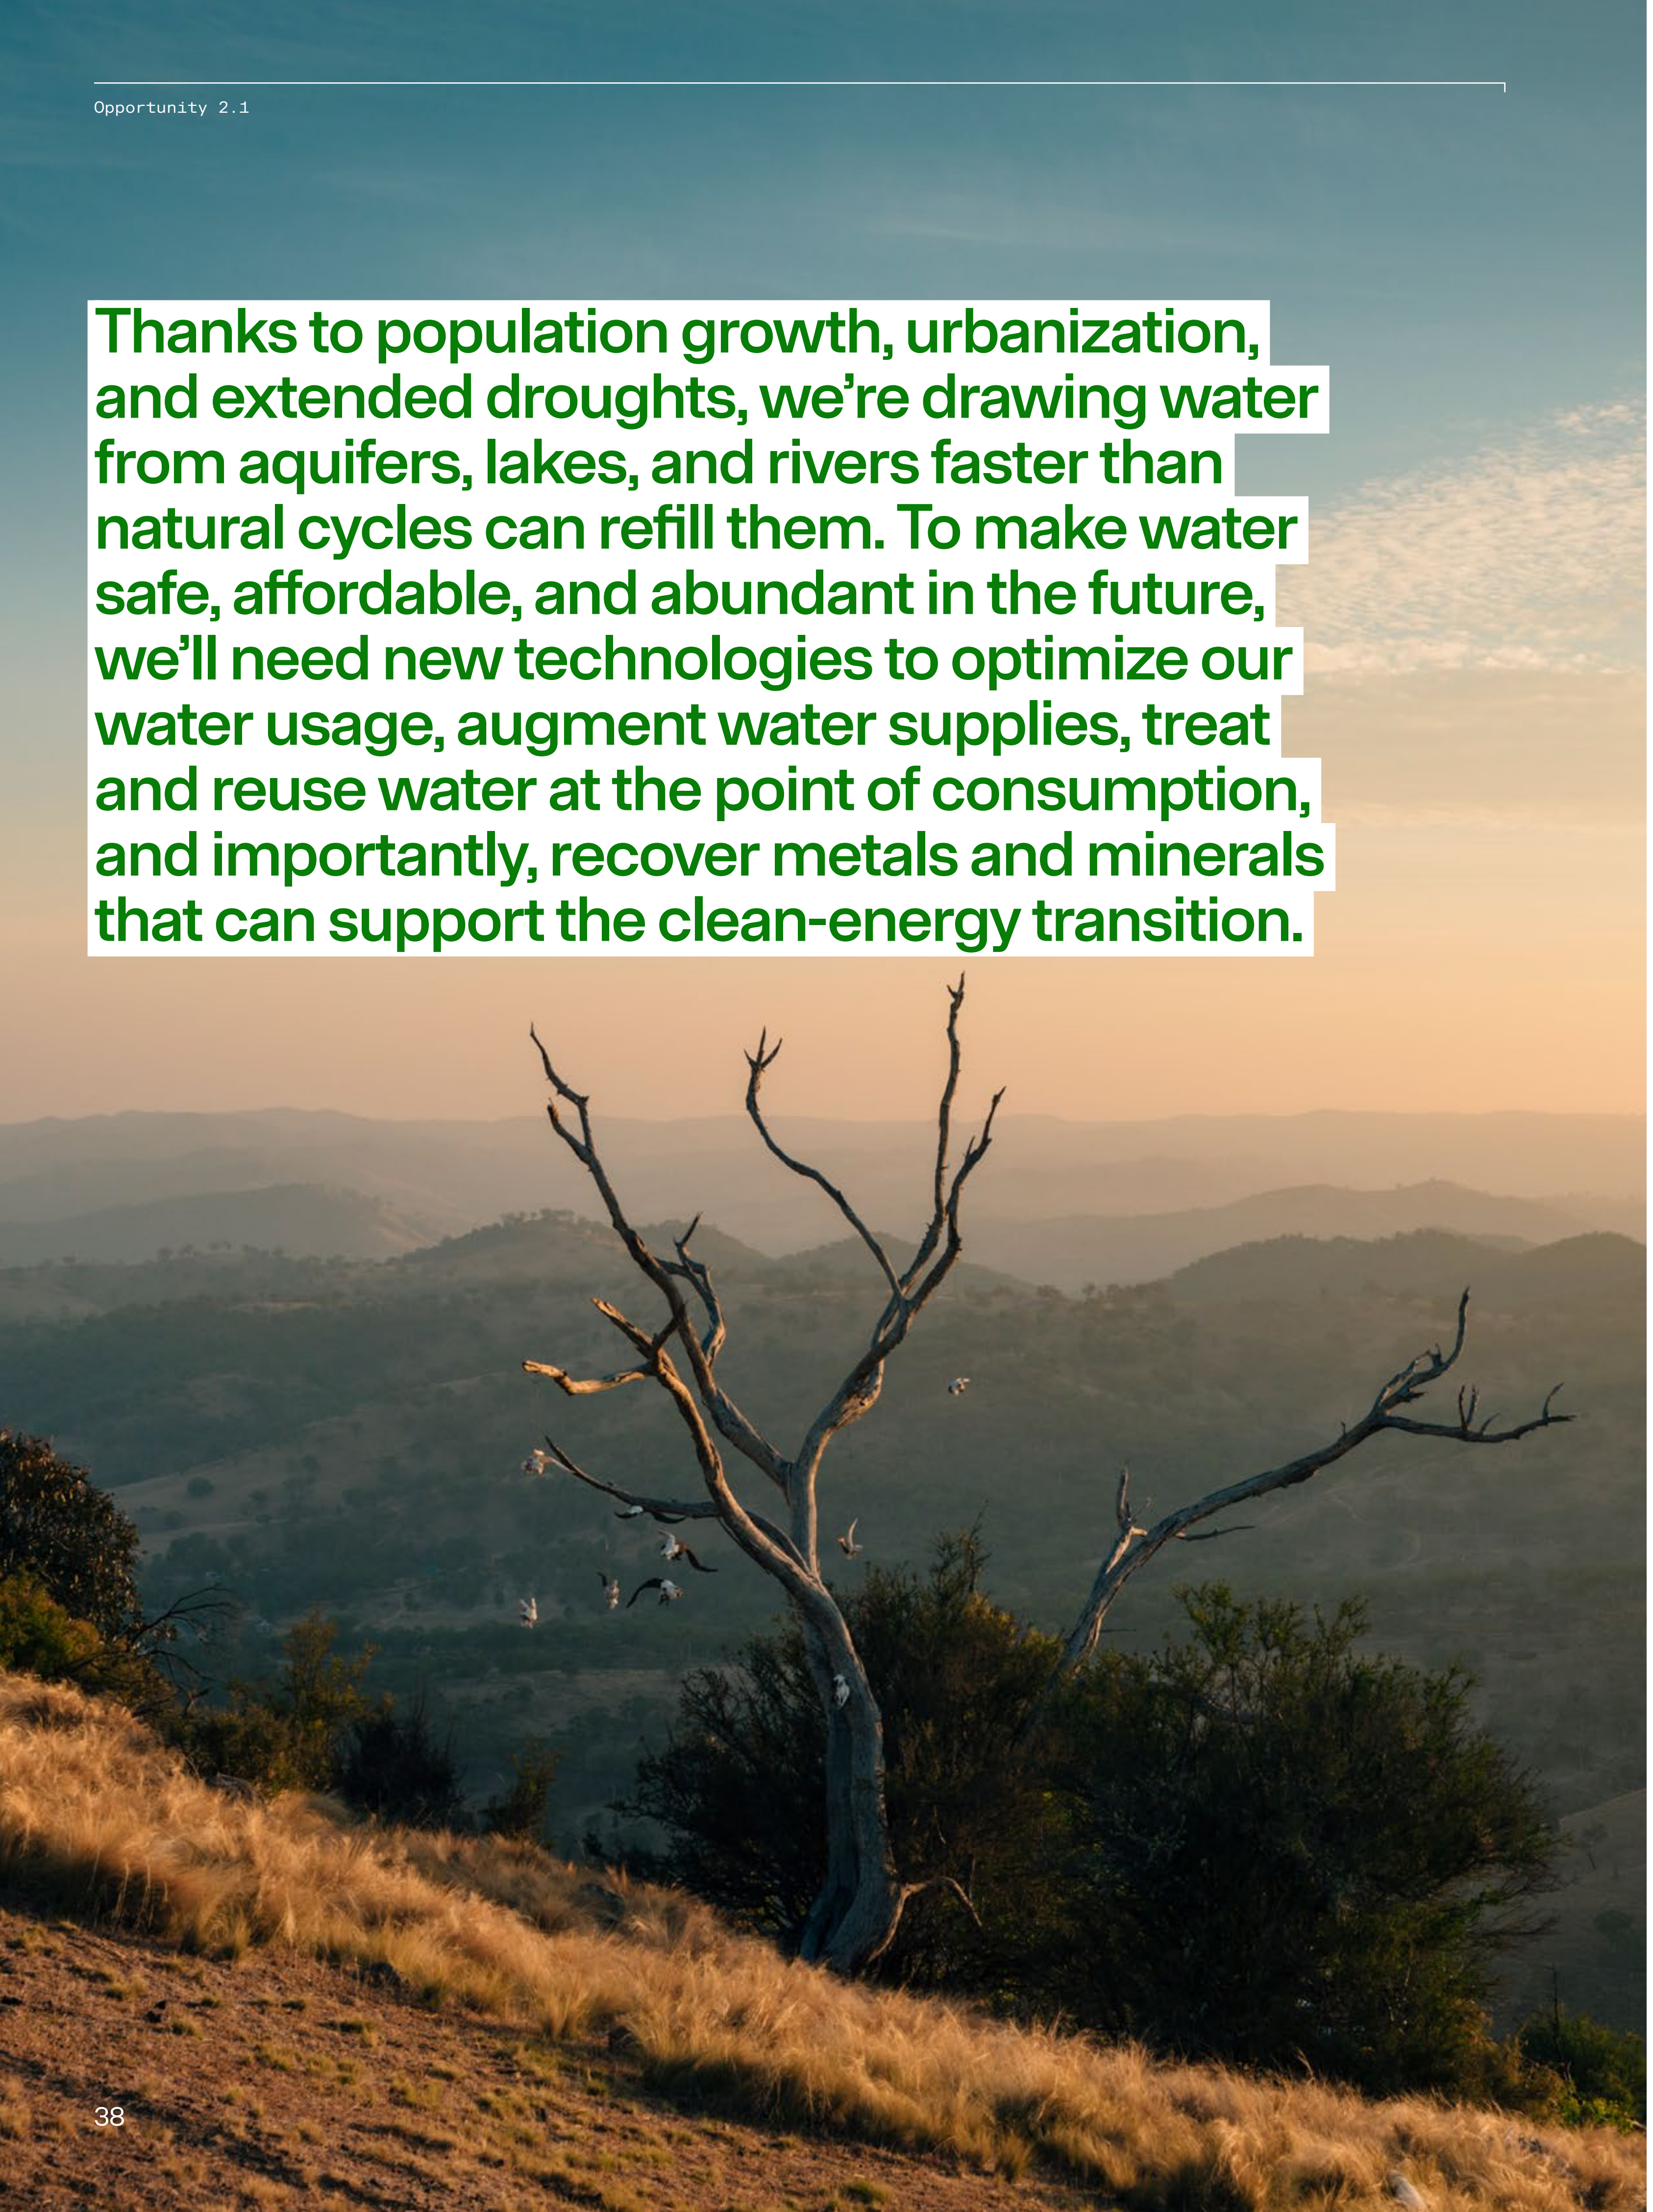

Thanks to population growth, urbanization, and extended droughts, we're drawing water from aquifers, lakes, and rivers faster than natural cycles can refill them. To make water safe, affordable, and abundant in the future, we'll need new technologies to optimize our water usage, augment water supplies, treat and reuse water at the point of consumption, and importantly, recover metals and minerals that can support the clean-energy transition.

By 2050, 4.8 to 5.7 billion people—more than half the global population—will live in areas where water is scarce at least one month per year. Climate change exacerbates the problem, causing protracted droughts, changing rainfall patterns, worsening flooding, and leaving thinner snowpacks each winter, among other problems. Jason Pontin, a general partner at DCVC, describes water shortages as the "shark's teeth" of climate change—the effect most people will feel first. "Long before we have two degrees of warming, water will be the singular challenge for our economies and security," he says, "unless we take dramatic action, now."

Yet despite the scale of the industry, water is one of the most neglected sectors in venture capital. Of the $40 billion investors put into climate tech in 2021, only $400 million went toward water technology, according to veteran water investor Tom Ferguson from Burnt Island Capital.

There are certainly reasons for the lack of venture investment in water. It's a highly regulated industry where buyers tend to be risk-averse and slow to adopt innovations. Water is also heavily subsidized; inelastic prices make it harder to get returns on innovation. And while there are many attractive opportunities for using AI to optimize water management, water is ultimately a physical asset and must be physically treated, which lengthens the time it takes to test and deploy new technology. The time- and capital-intensity of the water business means fewer companies soak up more money, a framework that spooks many investors.

All of that puts a premium on deep domain experience. "Water issues are here now and acute," says Earl Jones, a DCVC operating partner and globally recognized expert on water. "Yesterday's solutions cannot meet our challenges of today and of the decades to come." Learning to treat and recycle water at the point of use will be one of the key ways out of the crisis. The practice is rare outside of water-scarce regions; Israel reportedly purifies nearly 90 percent of its household wastewater and uses it for agricultural irrigation, but the United States reuses less than 10 percent. There's little reason to fix issues such as single-use water, flood irrigation, and a leaky water infrastructure when governments around the world subsidize the cost of water and wastewater services to citizens, companies, and farmers—making water so cheap in many places that people don't have to track their monthly bills, or can't without purchasing their own equipment.

But one company seeking to change the economics of water reuse through deep tech innovation is ZwitterCo, a DCVC portfolio company based in Boston. Industrial water users would recapture more of their wastewater through membrane filtration if they could; the problem is that wastewater, by its nature, is full of contaminants that rapidly foul membranes. The chemicals typically used to remove these contaminants are so harsh that the membranes themselves quickly degrade and must be replaced—sometimes every few days.

ZwitterCo has built on research at Tufts University to develop a new membrane chemistry using zwitterionic polymers. The membranes can repel the fats, oils, and greases found in many types of wastewater as well as chlorine, the most common disinfectant. Most important, they recover easily from fouling. Gentle chemicals can restore ZwitterCo membranes to as-new condition over and over again.

That makes membrane filtration a viable option for wastewater reuse at facilities like dairies, which have high organic contamination in their wastewater that's nearly impossible to remediate using traditional membranes. Pharmaceutical companies use ZwitterCo technology not only to recover water, but also to concentrate and recover high-value proteins. "The ticket that really gets people excited," says ZwitterCo cofounder and CEO Alex Rappaport, "is when you help them exceed their sustainability goals *and* create unexpected revenue opportunities." He says the company expects to extend its technology to all categories and forms of water filters.

▸ Irrigation for agriculture accounts for 70 percent of freshwater consumption worldwide and 80 percent in the United States, according to the OECD and the USDA.

> **"The risks posed by our global water situation are clear: headwinds to economic growth, negative impact on population health, increased risk of catalyzing regional conflict, and acute risk to food security.**

Earl Jones

Operating Partner, DCVC

ZwitterCo isn't the only company developing techniques for recovering valuable materials from wastewater. The brines from geothermal energy, mining operations, oil and gas extraction, and seawater desalination contain a host of metals and minerals, including lithium, manganese, cobalt, phosphorus, potassium, and magnesium, as well as calcium and sodium salts; recovering these resources could yield materials for batteries, solar panels, and other parts of the renewable-energy infrastructure. Jones calls this approach "wastewater mining" and says it can "help move us from a resource-disposal economy to a pervasive resource-recovery economy."

▸ **ZwitterCo** has built on research at Tufts University to develop a new membrane chemistry using zwitterionic polymers.

That's the logic behind DCVC's investment in Auckland-based Aquafortus. The company rejects wasteful, ecologically damaging traditional wastewater-mining approaches such as evaporation ponds, like those used at Salar de Atacama in Chile for lithium recovery, or other technologies that involve complex combinations of chemical precipitation, membrane concentration, thermal evaporation, and ion exchange. Aquafortus' solvent-extraction process starts by mixing a novel absorbent into a wastewater stream. The absorbent pulls water from the wastewater, allowing metals and minerals to precipitate out. The final step of the process separates the water from the absorbent, yielding fresh water and absorbent that can be used again. Aquafortus' solution reduces the energy cost of wastewater mining by up to 90 percent compared to evaporation ponds and chemical precipitation, and reduces the overall cost by up to 60 percent.

Even as we find ways to reuse wastewater and its contents, we'll need to deal with the toxins that contaminate our remaining water supplies. Per- and polyfluoroalkyl substances (PFAS), used for over 70 years as the raw materials for industrial products like Teflon and fire-retardant fabrics, are a particular problem: they don't degrade naturally, earning them the nickname "forever chemicals," and they've become pervasive in our soils, waterways, and bodies. The U.S., Canada, Europe, and Australia are aggressively regulating PFAS and the multidecade process of environmental PFAS remediation has begun.

One of the companies working on solving the problem of PFAS contamination in water is Boston-based Aclarity, whose seed funding was led by DCVC in 2022. Founder Julie Bliss Mullen invented an advanced electro-oxidation solution that destroys PFAS compounds in water and wastewater. "Aclarity has proven its ability to destroy PFAS," Jones says. "But what makes Aclarity such a great investment opportunity is its ability to do that using low energy, which correlates to low cost, as well as its outstanding founding team." In a full-scale pilot project with the $19 billion water technology firm Xylem, Aclarity proved that it could reduce PFAS levels in landfill leachate from more than 1,000 nanograms per liter to less than 10 nanograms per liter, using very little energy.

Aclarity, ZwitterCo, and Aquafortus are platform companies that bring industry experience to the chemistry, electrochemistry, and materials science of water treatment. Each targets large markets, using ideas that outflank existing solutions both technologically *and* economically; DCVC's outlook is that large-scale adoption of sustainable solutions in the water industry can only come through dramatic reduction of both capital expenditures and operating expenses.

"The risks posed by our global water situation are clear: headwinds to economic growth, negative impact on population health, increased risk of catalyzing regional conflict, and acute risk to food security," Jones says. "Smart investors see the challenge and see the opportunity and want to put their money to work and be part of solving water. At DCVC, we are making an impact on global water through innovation and we're going to make money while we're doing it."

Figure 2.1 Stricter Standards

Tests by the Environmental Working Group have found PFAS levels above 70 nanograms per liter (also expressed as parts per trillion) in the tap water in two U.S. locations (Brunswick County, NC, and Quad Cities, IA) and between 4 and 60 ng/l in 30 more locations. Between 2016 and 2023 the Environmental Protection Agency drastically lowered its guidance for the amount of perfluorooctanoic acid (PFOA) and perfluorooctane sulfonic acid (PFOS) that is allowable in drinking water. (Both are members of the PFAS chemical group.) The Infrastructure Investment and Jobs Act of 2021 allocates $5 billion to addressing PFAS and other emerging contaminants in drinking water.

	Pre-2016 (ng/l)	2016 EPA Health Advisories (ng/l)	2023 EPA Proposed Standards (ng/l)
PFOS	200	70*	4.0
PFOA	400		4.0

* Combined contaminants

Aclarity proved that it could reduce PFAS levels in landfill leachate from more than 1,000 nanograms per liter to less than 10 nanograms per liter, using very little energy.

▶ The transition from fossil fuels to electric power will accelerate as we identify simple, elegant solutions for many of the remaining problems.

Featured Companies

Agora Energy Technologies
Antora Energy
Form Energy
Reach Labs [DCVC]
Verdigris [DCVC]

Climate Tech: Electrify Everything

In every case where carbon can be taken out of the energy equation, it must be. That means replacing almost every fossil-fuel-burning machine in industry, transportation, and building management with electric alternatives, and decarbonizing and shoring up the grid to power them.

As a form of energy, renewable electricity is simply better than fossil fuels. For one thing, electrical energy can be directly transformed into light, heat, or motion—whereas the chemical energy in hydrocarbons, before it can become equally versatile, must be converted into thermal energy (via pollution-creating combustion), then converted into kinetic energy (via a steam cycle), and only then converted into electrical energy (by a generator). In most cases, electricity is also easier to transport, and it generates no greenhouse gas emissions at the point of use.

Those facts give us a huge head start as we contemplate electrifying the entire economy. Now we just need to *do* it. We must stop selling new fossil-fuel-burning equipment, then decommission every internal-combustion automobile engine, every gas-burning furnace and kiln, every diesel-burning truck and cargo vessel, and every kerosene-burning regional passenger plane and substitute modern electric-powered versions.

We're just at the beginning of this transformation. The revolution will require a combination of public-policy incentives, private investment, and technical innovation, especially for the infrastructure to deliver electricity on demand to every machine that will need it. DCVC's portfolio in this area is still small, but our investing professionals are on the hunt for companies with demonstrated success and a path to scale. "We're looking for solutions that are simple and elegant, and we haven't found all the ones that we want yet," says DCVC partner Rachel Slaybaugh.

Figure 2.2 Projected EV Stock, United States

There will be more than 18 million battery-powered electric vehicles and plug-in hybrid electric vehicles on U.S. roads by 2030, the International Energy Agency projects.

BEV PHEV

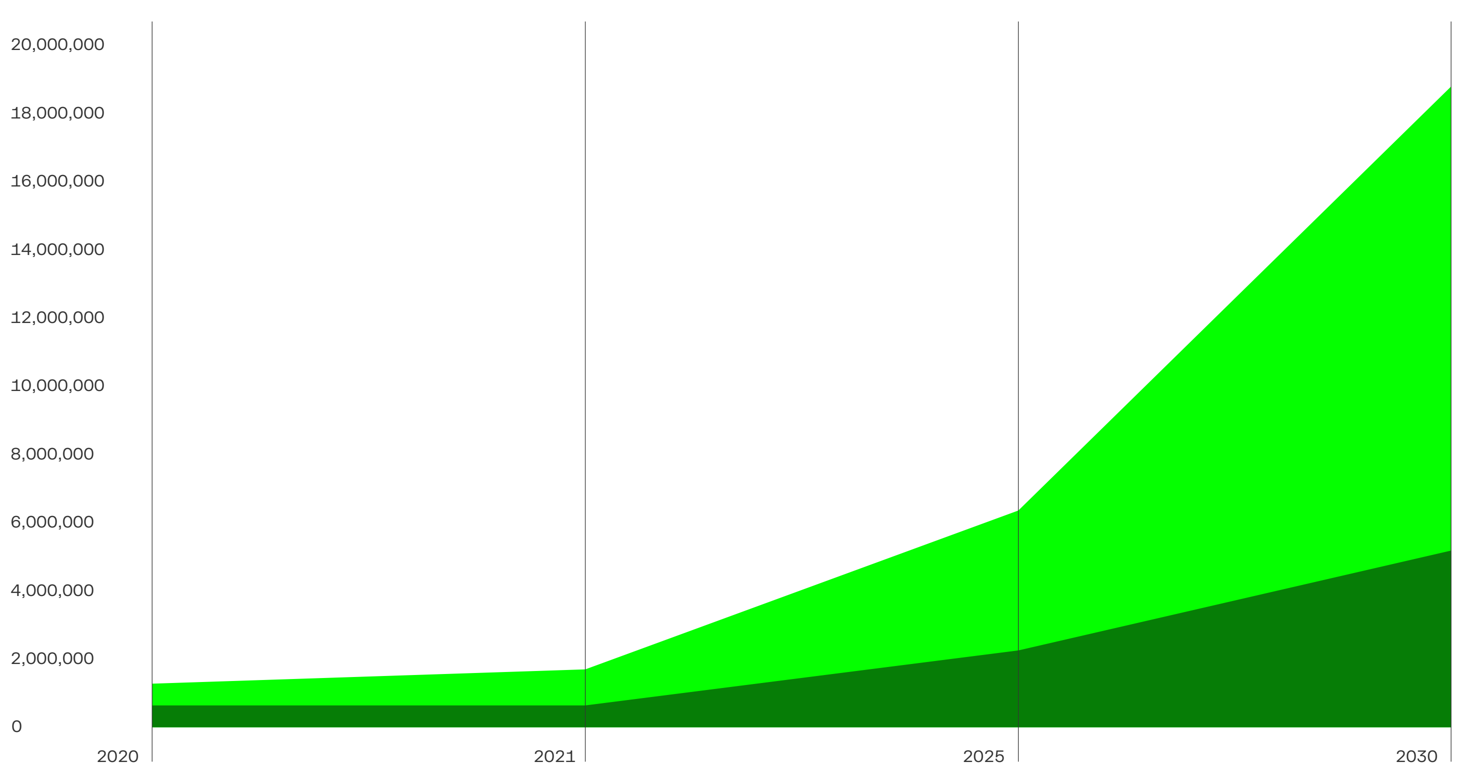

One area we see as ripe for innovation is home energy storage. Equipping all homes so that they have the power supply needed to charge future electric vehicles—as well as other electrical devices, such as the heat pumps that will replace furnaces and air conditioners—will raise system-level challenges. If everyone in a neighborhood installed chargers for their new EVs, for example, it would likely overwhelm local power lines and transformers. Companies like Agora Energy Technologies, Antora Energy, and Form Energy are exploring new battery chemistries and form factors for microgrid-scale storage—but what we haven't invested in yet is a company building a home energy-storage appliance to cover brief periods of intense electrical consumption.

What's needed is a new category of power-storage appliances, available for $5,000 to $10,000, that could be installed in any residential or commercial building to provide electricity in short bursts of 5 to 10 minutes, at times of day when the grid is overtaxed and the price per kilowatt-hour is high. Such appliances could recharge when electricity rates are lower. Here at DCVC, we have our eyes open for the team that figures this out first.

Another elegant way to retire combustion-powered machines and shift more work to electrical devices would be to untether them from the grid by sending power through the air. That's what DCVC-backed Reach is working to do. Its AI-based power-at-a-distance technology can steer radio-frequency electromagnetic energy to a mesh of devices in factories, warehouses, and stores. (It's like Wi-Fi, but at higher power levels, with safety protocols that protect living things by disabling or rerouting the power before anything enters the beam.) Reach's system can already power small devices such as displays, video cameras, industrial sensors, and inventory trackers, and as the technology matures it will power robots, autonomous guided vehicles (AGVs), drones, and entire buildings.

"If you look out your window at those power lines, you're receiving your electricity pretty much the same you have since Thomas Edison and George Westinghouse were slugging it out in 1910," says Alan Cohen, a DCVC general partner and Reach board member. Wireless power-at-a-distance gets electricity to places where it's not efficient, or not possible, to wire things up physically or to repeatedly recharge their batteries, Cohen says. "DCVC very frequently uses deep tech to go after real-world industrial problems and opportunities, and the ability to transmit power to this vast array of targets is an absolutely classic case. Reach's opportunity is to become the leader of the wireless microgrid."

Preparing for a future where millions of people drive EVs and millions of homes have energy-producing rooftop solar panels will require a massive upgrade of the electrical grid.

DCVC-backed companies continuing to chip away at other important parts of the electricity delivery and management problem include Petra, which is testing a new way of tunneling through rock that would drastically lower the cost of installing underground transmission lines, and Verdigris, which makes an "energy data gateway" that attaches to a circuit-breaker panel to help analyze a building's electricity use.

To win large-scale adoption, any new technology intended to advance electrification will need to be user-friendly, clicking into rather than clashing with existing infrastructure. "Systems can become really complicated and hard to use, and people just get bogged down in permitting and installation," Slaybaugh says. "We need solutions that can actually be adopted in practice. Not that many people are thinking this all the way through yet."

▸ Innovation in nuclear power, both fission and fusion, is experiencing a renaissance thanks to advanced computing and manufacturing technology.

Featured Companies

Radiant Nuclear [DCVC]
Zap Energy [DCVC]
Oklo [DCVC]
Commonwealth Fusion Systems

Climate Tech: Next-generation Nuclear Power

The United States and other nations must vastly increase their capacity to generate dispatchable, carbon-free electricity. Advanced nuclear fission is a real and increasingly practical alternative to burning coal or natural gas, thanks in part to advanced computation that's leading to safer, smaller, easier-to-build plant designs. And fusion energy, given the recent announcement of scientific energy breakeven at a U.S. Department of Energy lab, is closer to commercialization than it's ever been.

Oklo's microreactors will produce 1.5 to 15 megawatts of electricity and can be built near customer sites.

Limiting global warming to manageable levels, while still leaving room for population and economic growth, will require so much additional clean energy capacity that we'll need many options. Two areas where DCVC is investing are advanced nuclear fission and nuclear fusion, which—unlike solar and wind—promise reliable electric power, as needed, at any time of day, in any season. New materials, enhanced sensors, and especially the ability to create sophisticated computational models of reactors and control systems are helping companies solve basic science and manufacturing problems that have stymied the traditional fission power industry for decades.

Our investing lens in this area is shaped by the deep experience of our team members. Slaybaugh, a partner at the firm, spent eight years as a professor of nuclear engineering at the University of California, Berkeley, while simultaneously starting the nuclear fission program at ARPA-E, the Department of Energy agency that funds early-stage energy research. She founded the Nuclear Innovation Bootcamp, an annual event that mentors students and early-career professionals, and in 2020–21 she ran the nuclear portion of the Department of Energy agency review on the Biden-Harris transition team. "I have worked with national laboratories, government agencies, and companies in the nuclear sector that help to train the next generation and help to find the problems that need to be solved for nuclear to be viable," she says.

Those problems will be a lot easier to solve with the aid of the latest computing technology. "The nuclear power industry had a lot of good ideas in the 1950s and 1960s, but then ran into implementation problems in the 1970s and 1980s," says Zachary Bogue, DCVC managing partner. "And for the next twenty years nobody really worked on solutions. But now that we finally have powerful simulation capabilities, you can revisit all the roadblocks and think about them differently this time, because computation can inform better materials and better processes." The first generations of U.S. fission power plants were built at ungodly expense mainly due to site-specific design and construction practices. Slaybaugh and Bogue believe these plants can now be replaced by cheaper, more flexible designs.

Radiant Nuclear is a DCVC-backed startup where a team of former SpaceX employees is using extensive computer modeling and digital twinning to rethink old-fashioned fission. A traditional fission reactor is several stories high and requires an elaborate cooling system to prevent its uranium fuel from melting down. But Radiant's microreactor is small enough to fit in a semitrailer. It's fueled by ceramic-coated uranium particles that are cooled by helium gas and can't melt down. In operation, the machine could generate as much electricity as a trailer-mounted diesel generator—that is, enough to power a military installation, medical clinic, or small town. Radiant plans to conduct a fueled test of its microreactor in 2026.

Oklo, another DCVC-backed startup designing a fission microreactor, is also using digital twinning and machine learning to improve the efficiency of fuel production for its proposed fast neutron reactor, a design studied for decades but never commercialized in the U.S. Like Radiant's microreactor, Oklo's demonstration reactor, called Aurora, would use a fuel called high-assay, low-enriched uranium (HALEU), which contains more of the fissile isotope U-235 than traditional nuclear fuel, allowing Aurora to provide more power per volume than the 93 light-water reactors currently operating in the U.S. The company is working with Idaho National Laboratory and Argonne National Laboratory to create the HALEU it needs by recycling nuclear waste from older reactors.

Passive safety features mean Oklo's design is also meltdown-proof. And because they're smaller than traditional reactors, requiring less capital investment, Oklo plans to own and operate its own plants and sell power to customers like universities, hospitals, and data centers. "All you need to do is sign a power purchase agreement with us, and we'll take care of the licensing, construction, refueling, and eventually decommissioning of the reactors," says Bonita Chester, Oklo's director of communications. "We can site where power is needed and allow communities to build resilience."

New computational techniques are having an even more dramatic impact in fusion research. Commonwealth Fusion Systems (CFS), in Massachusetts, isn't backed by DCVC, but its story is an inspiration to nuclear entrepreneurs everywhere.

The first milestone in the race to build a practical fusion plant was to get more energy by fusing atomic nuclei than it took to force the nuclei together. Researchers at the Department of Energy's National Ignition Facility in Livermore, California, achieved that goal in December 2022, briefly obtaining 3.15 megajoules of energy in an inertial confinement fusion experiment that required only 2.05 megajoules of input.

But while it was an exciting proof of concept, the NIF style of fusion—in which a stadium-sized array of lasers is used to compress a fuel capsule until it implodes—would take decades to implement at any practical scale. Tokamak-based designs, in which powerful magnets are used to confine a doughnut-shaped ring of plasma, are far closer to the commercial stage. A central challenge in tokamakstyle fusion, though, has been designing magnets powerful enough to confine the plasma, but still compact and cheap enough to be manufactured at scale. That's the focus of CFS. The company figured out how to use a novel high-temperature semiconductor material called rare-earth barium copper oxide (REBCO) to build smaller, cheaper magnets.

CFS demonstrated in 2021 that its prototype magnet can create a 20-tesla field—enough to sustain fusion in the company's planned demonstration reactor. And that innovation was greatly accelerated by the development of a high-fidelity digital twin of the magnet. "You build a physical system, informed by your model. You instrument that physical system and get tons of data that make the model better. Then you basically run them in parallel until the digital twin is a faithful mirror of reality," Bogue explains. "That's how you iterate more quickly on the fundamental innovation, and that's exactly how CFS built their magnets."

Zap Energy, a DCVC portfolio company in Washington state, is closing in on yet another type of fusion machine, but with equal help from simulations. When an electric current runs through a metal tube such as a hollow lightning rod, it creates a magnetic field around and inside the tube—and if the current is strong enough, the resulting magnetic field can crush the tube. That's called the Z-pinch effect. Zap Energy is building a prototype fusion reactor that works the same way, but in place of the metal tube, it energizes deuterium-tritium gas to form a two-foot filament of plasma with its own magnetic field.

To suppress instabilities and keep the filament from dissipating, the company exploits a phenomenon called sheared flow stabilization; it's akin to creating a multilane freeway with traffic moving at different speeds in different lanes. And as it builds its prototype fusion core, Zap is using simulations to model all the equipment and techniques needed to achieve flow stabilization.

As we noted in the introduction to this report, the company generated its first plasmas in 2022. Their ultimate goal is to build fusion cores that don't need superconducting magnets to confine the plasma and are therefore even more compact and less expensive than CFS-style tokamaks.

For Bogue and Slaybaugh, it's natural to expect that new ideas in nuclear fusion and fission will come from disciplines like computer science and aerospace. The founders of Radiant, in fact, are the same people who figured out how to make SpaceX's Falcon 9 rocket land upright. "Building the control system for that was a digital-twin exercise," Slaybaugh observes. "Now they're applying that engineering expertise to fission—and they say it looks like an easier problem."

> "You build a physical system, informed by your model. You instrument that physical system and get tons of data that make the model better. Then you basically run them in parallel until the digital twin is a faithful mirror of reality.
>
> **Zachary Bogue**
> Managing Partner, DCVC

Figure 2.3.1 Annual Greenhouse Gas Emissions Abated by the U.S. Nuclear Industry

If the power entering the U.S. energy grid from nuclear plants had come from fossil fuel plants instead, CO_2 emissions would have been 500 to 750 megatons higher each year since 1995; in total, nuclear plants have kept almost 17 gigatons of CO_2 out of the atmosphere over that time. Unfortunately, the amount of emissions avoided each year is declining as conventional nuclear reactors go offline.

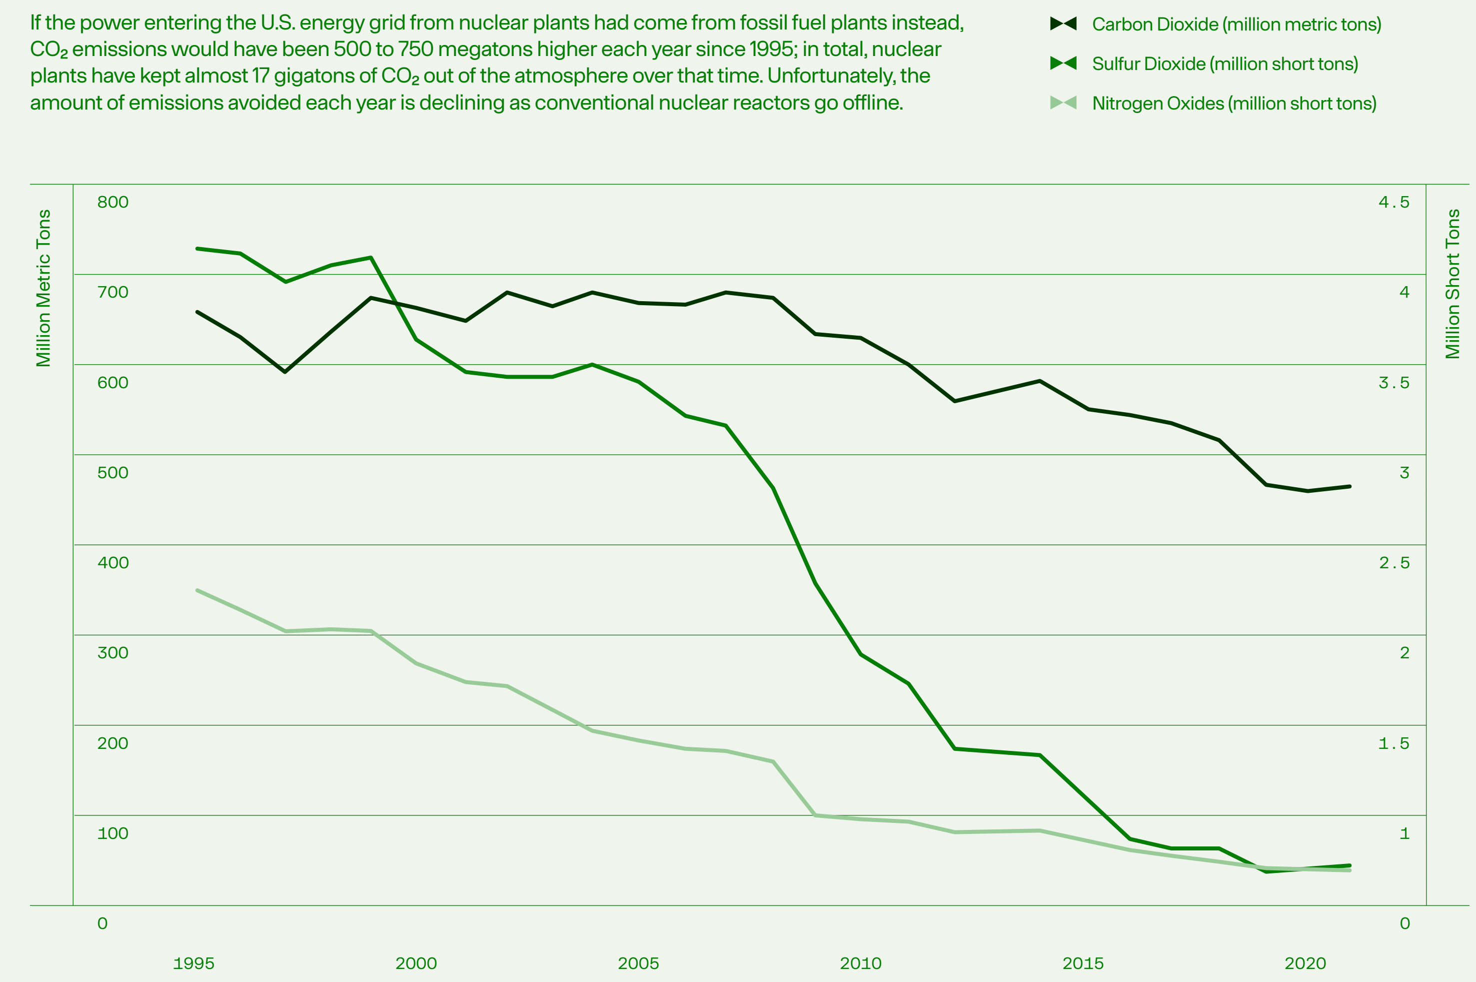

Source: Nuclear Energy Institute; U.S. Environmental Protection Agency; U.S. Energy Information Administration

Figure 2.3.2 Construction Delays

In Western Europe and the United States, three traditional light-water nuclear reactors are under construction or recently completed. All were afflicted by massive cost and schedule overruns.

Project	Construction Start Date	Original Projected Startup Date	Projected Actual Startup Date	Original Projected Cost	Current Estimated Cost
Vogtle Electric Generating Plant Georgia	2012	2017	2023	$14 B	$30 B
Flamanville-3 Nuclear Plant France	2007	2012	2023	$3.5 B	$14 B
Hinkley Point C Nuclear Plant United Kingdom	2017	2013	2028	$19 B	$31 B

Direct Air Carbon Capture and Storage

A Climeworks facility near the Hellisheidi power plant outside Reykjavik, Iceland

In case you've run out of other things to worry about, put this on your list: One of the main technologies we'll need to arrest global warming—direct air capture of CO_2—is economically out of reach. At least for now.

↳ In a 2018 special report, the U.N.'s Intergovernmental Panel on Climate Change said that to limit global warming to 1.5°C above preindustrial levels, it won't be enough to stop adding greenhouse gases to the atmosphere. We'll also need to *remove* somewhere between 100 and 1,000 gigatons of CO_2 from the air before the century is out. "Unless affordable and environmentally and socially acceptable CDR [carbon dioxide removal] becomes feasible and available at scale well before 2050," the panel predicted, we'll significantly overshoot the 1.5°C goal.

The main types of CO_2 removal the panel examined were burning biomass fuels with full carbon capture and storage; reversing deforestation and planting trees in areas not previously forested; sequestering carbon in the soil through regenerative farming methods; enhanced weathering, which involves spreading finely ground silicate rocks on land or in the oceans; and direct air carbon capture and storage, which requires separating out molecules of CO_2 that are present in the atmosphere at about 417 parts per million. The panel judged that all of these methods face "multiple feasibility constraints." The main problem with direct air capture, it said, was that it was too energy-intensive, and therefore too expensive.

That will still likely be the case in five, seven, even 10 years—which is why we at DCVC are somewhat surprised to see hundreds of millions of dollars in capital flowing into early-stage direct air capture companies, such as Switzerland's Climeworks, Canada's Carbon Engineering, and U.S.-based Global Thermostat.

These companies are building small demonstration plants where air passes over a solid absorbent filter or a liquid solvent that takes up CO_2 molecules. The filter or the solution is then heated to 80–900°C to release the CO_2, which is pumped underground or used for industrial purposes. The heating phase requires a lot of energy, and so far, these companies haven't shown how they'll bring their energy needs down enough to compete with other forms of carbon removal.

A Climeworks facility near Reykjavik charges customers a reported $600 per ton for its carbon offsets (even though the plant's own power comes from geothermal energy). Carbon Engineering says its technology can capture CO_2 at a cost of $100 to $250 per ton. But do the math. Even at $100 per ton, capturing 100 gigatons of CO_2, the *low* end of the IPCC's estimate of what will be necessary, would cost $10 trillion—more than a tenth of global GDP.

Obviously, the technology will advance, especially now that the federal government is investing in the technology. The 2020 federal stimulus bill gave the Department of Energy $447 million to set up grants, competitions, and testing centers for air capture technology. And one geothermal developer backed by DCVC, Fervo Energy, is taking advantage of a grant from the Chan Zuckerberg Initiative to build its own direct air capture plant, as a way to test the economics of powering solid-sorbent direct air capture with waste heat from geothermal wells. The question for careful investors is whether the needed progress on the core technology behind direct air capture will come soon enough to pay off the earliest bets. On the online prediction marketplace Metaculus, the median guess for when direct air capture costs will fall below $50 per ton is 2040. We think that's optimistic.

Point-source carbon capture is evolving much faster. We're watching companies like Carbon Clean, which builds modular chambers that use a liquid solvent to absorb flue gas at cement and steel plants and oil refineries; Remora, which can capture at least 80 percent of the CO_2 coming out the tailpipes of semitrucks; and Osmoses, which is developing novel gas-separation membranes that could lower the cost of carbon capture.

Such technologies could become affordable in the foreseeable future, and we hope to see them in operation someday on the roads and at every industrial site that emits CO_2. But in Slaybaugh's estimation, it's "just too soon" for direct air capture to be a viable climate-tech investment opportunity.

"I'm not saying we won't need it," Slaybaugh says. "And I'm not saying there won't eventually be good businesses here. I'm saying right now the markets are very nascent, and I don't see how you can possibly make a venture return."

Hydrogen-fueled cars, buses, trucks, and generators release no greenhouse gases as they work, leading some to tout hydrogen as a clean alternative to fossil fuels. But hydrogen lacks the energy density of petroleum-based fuels. It's also energy-intensive to produce while being difficult and expensive to store and transport safely. We don't think the compromises currently required are worthwhile. We believe hydrogen is, in part, a pathway touted by the petrochemical industry because it furthers its existing natural gas businesses.

Figure 2.5	The economics and the environmental benignity of different colors of hydrogen

Hydrogen fuels are assigned colors (gray, blue, sage, green, lime, and yellow, in this schema) depending on what feedstocks are used and how production is powered. Only green and lime hydrogen are carbon-free. SMR = steam methane reformation, CCS = carbon capture and storage

Primary Inputs	Natural gas + water, coal + air	Natural gas + water + air	Natural gas + water	Renewable electricity + water	Nuclear electricity + water	Grid electricity + water
Process	SMR, coal gasification	SMR + CCS	Pyrolysis	Electrolysis	Electrolysis	Electrolysis
Output	$H_2 + CO_2$	$H_2 + CO_2$ (partially captured)	$H_2 + C$	$H_2 + O_2$	$H_2 + O_2$ + nuclear waste	$H_2 + CO_2$

Source: A. Ajanovic, M. Sayer, R. Haas, "The economics and the environmental benignity of different colors of hydrogen," International Journal of Hydrogen Energy (Vol. 45, Issue 57, 5 July 2022, pp 24136–24154)

↳ A wave of enthusiasm for hydrogen hits every 15 years or so, and we're in another one right now, thanks to huge new tax credits for "green hydrogen" production and investment built into the U.S. Inflation Reduction Act of 2022.[1] "But it hasn't yet proven to be a breakout idea," Bogue argues. "It's an inferior-performing fuel. But that hasn't stopped the incumbent petrochemical industry from embracing it, and that should tell you something. Still, we are open to seeing the kind of innovations that would make this idea more promising."

When H_2 gas is burned, the only outputs are heat and water, which means hydrogen is in some sense clean. But making hydrogen certainly is not. About three-quarters of the world's H_2 gas is produced through steam reforming of natural gas (creating "gray hydrogen"), which releases 10 kilograms of CO_2 for every kilogram of H_2 produced ($CH_4 + 2H_2O \Rightarrow CO_2 + 4H_2$). Making green hydrogen involves a different process: separating water into oxygen and hydrogen through electrolysis, ideally using electricity from renewable sources. But electrolysis is two to four times as expensive as steam reforming, and today only 1–4 percent of hydrogen is produced this way.

As an energy carrier, meanwhile, even green hydrogen is unimpressive. An analysis published by the World Economic Forum found that "approximately 30–35 percent of the energy used to produce hydrogen is lost during the electrolysis process; liquefying or converting hydrogen to other carriers, such as ammonia, results in a 13–25 percent energy loss; and transporting hydrogen requires additional energy inputs that are typically equal to 10–12 percent of the hydrogen's own energy." By the time hydrogen fuel reaches its end user, in other words, only about 25 percent of the energy needed to extract, package, and move it is left over for practical use.

That's just too inefficient to make economic sense at present, outside of a few key applications. That said, nearly all promising pathways to decarbonize the heavy transport industries (air travel, shipping, long-haul trucking) involve burning H_2 directly, or as one of the inputs in creating sustainable versions of existing fossil fuels like "SAF" (sustainable aviation fuel). And green hydrogen in fuel cells looks like a reasonable option for applications that electricity from renewable sources can't reach, such as spacecraft or submarines. (In fact, DCVC backs a startup called Amogy whose ultra efficient catalyst converts ammonia, NH_3, into hydrogen that can be fed directly into fuel cells.) Ergo, a specialized hydrogen industry could emerge thanks to the new tax credits. But as a one-to-one substitute for fossil fuels, hydrogen is mediocre on every level.

So why are big petroleum companies like Shell and BP championing hydrogen? Possibly because most of it is made from one of their core products, natural gas. Their refineries also need a lot of hydrogen for processes such as hydrocracking—breaking the hydrocarbons in crude oil into simpler molecules like gasoline and kerosene—and removing sulfur. But there may be another layer to their strategy. As long as today's big fossil fuel consumers see hydrogen as a real option, they may postpone the switch to cleaner, fossil-free energy sources such as solar, wind, geothermal, and nuclear power.

The International Energy Agency argues that while there have been false starts for hydrogen in the past, "this time could be different," now that governments, automakers, utilities, oil and gas companies, and municipalities all seem to be interested in the technology. Perhaps so. As investors, we're driven by evidence above all else, and we continue to monitor this area closely.

[1] Grey, blue, green – why are there so many colours of hydrogen? https://www.weforum.org/agenda/2021/07/clean-energy-green-hydrogen/

Carbon-intelligent Materials

We'll never fully eliminate carbon from the global economy, but we can get a lot smarter about how we use it.

Twelve x Pangaia

▸ **Twelve** worked with Pangaia to create the polycarbonate Lab Sunglasses, in which carbon captured from point industrial sources replaces fossil carbon.

If fossil fuel combustion were the only source of CO_2 emissions, switching to a carbon-free energy grid would be enough to slow global warming. Carbon, however, is an essential part of many irreplaceable materials, such as medical-grade plastic, and is also a byproduct of many common activities that can't be easily decarbonized, such as air travel and construction. DCVC enthusiastically supports companies developing less carbon-intensive and net-zero-carbon ways to make commodities such as cement, plastic, and jet fuel.

We applaud, for example, companies experimenting with electric-powered regional planes, but it's hard to imagine advances that would give batteries the capacity and low weight required for long-haul flights, like the 15-hour journey from Los Angeles to Sydney. The best we can hope to do in the near future is switch to green jet fuel that adds no net carbon to the atmosphere—meaning kerosene made with carbon diverted from industrial point sources.

That's exactly what Twelve—a DCVC-backed company in Berkeley, California—is making, along with a range of other net-zero-carbon materials. The founders began with the same electrolysis technology used to make hydrogen from water, then developed new electrode designs and new anode chemistry that let them run their electrolyzers at low temperature and high current density. The Twelve process takes in CO_2 from an industrial point source or direct air capture, humidifies it, and breaks the CO_2 and H_2O into CO, H_2, and O.

A CO-H_2 mixture is known as syngas, and it's the feedstock for a huge range of products, from plastic to green fuels. The materials science company Pangaia has used Twelve's syngas to make polycarbonate sunglass lenses, Procter & Gamble has tested it as a replacement for petroleum-derived materials in Tide detergent, and the U.S. Air Force helped Twelve test and certify a syngas-derived jet fuel called E-Jet.

If the electrolysis process used to convert the compounds is powered with clean energy, it adds no new carbon to the atmosphere, and depending on how the end product is used, it can actively lower atmospheric CO_2 levels. "With fuels, we can get to carbon-neutral, since you're going to combust it again as part of its use," says Nicholas Flanders, cofounder and CEO of Twelve. "With materials, you can potentially get to carbon-negative."

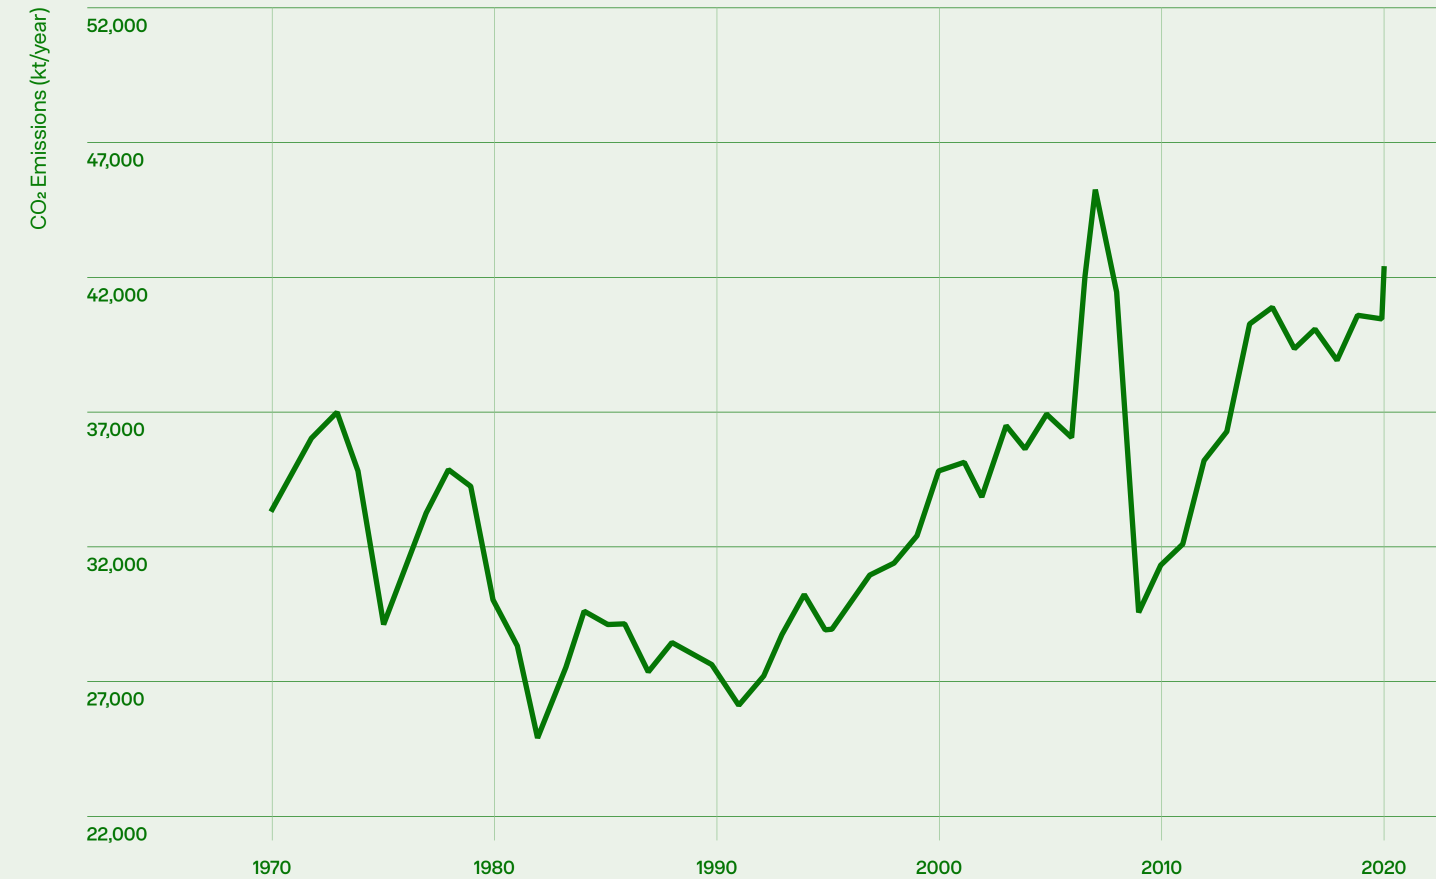

Figure 2.6
CO_2 Emissions From Cement Production: United States

About 8 percent of human-caused CO_2 emissions come from cement manufacturing, with the United States contributing roughly one-fortieth of that total. Although carbon emissions from U.S. cement production dropped after the 2008 financial crisis, they have since climbed back up to near-record levels.

> **"With fuels, we can get to carbon-neutral, since you're going to combust it again as part of its use … With materials, you can potentially get to carbon-negative.**
>
> **Nicholas Flanders**
> Cofounder and CEO, Twelve

Brimstone Energy makes cement from calcium-oxide-rich basalt rock rather than limestone.

Another DCVC-backed company working to create a carbon-negative version of a global commodity is Brimstone Energy. The company is rethinking cement production, which is responsible for a whopping 8 percent of human-caused CO_2 emissions. The traditional process for making ordinary portland cement (OPC), the world's most common building material, is to heat sedimentary limestone ($CaCO_3$) in large kilns until it breaks down into calcium oxide (CaO), or quicklime. Combined with other supplementary cementitious materials, or SCMs, such as slag from steel production and fly ash from coal burning, the quicklime forms cement powder. The problem is that the breakdown of limestone also releases a large amount of carbon, which combines with oxygen in the air to form CO_2—between 650 and 920 kilograms of it for every 1,000 kilograms of cement.

Brimstone sidesteps that fact of chemistry by eschewing limestone and starting with basalt and other igneous rocks, which are rich in native calcium oxide. Brimstone's process not only releases zero carbon, but it creates some useful byproducts, including the very SCMs needed to make OPC, as well as magnesium hydroxide and magnesium silicate hydrates, which actually absorb CO_2 from the atmosphere and sequester it in the form of magnesium carbonate. "Even if a Brimstone customer uses coal to fire their kiln, they're carbon-negative," says Slaybaugh. "And the cleaner the kiln firing, the more carbon-negative the cement."

We're also watching other companies that are working on making low-carbon versions of common materials. Sublime Systems is developing a way to make a type of calcium-silicate-hydrate cement using room-temperature electrochemistry instead of combustion-driven kilns. Boston Metal wants to shrink the vast carbon footprint of steel production using an innovative process called molten-oxide electrolysis. Traditionally, steel is made in a blast furnace where iron ore (Fe_2O_3) is combined with coking coal and limestone. Carbon released by the coal pulls oxygen away from the ore, which leaves behind highly pure iron but also creates huge amounts of CO_2—about 1.8 tons for every ton of steel produced. Boston Metal's process is far simpler: it uses electricity to heat iron ore until it releases high-quality liquid iron, with oxygen as the only byproduct.

The bigger point is that we don't need to fear carbon itself; it is, after all, the basis for all life on Earth. We just need to be smarter about how and when we use it. In particular, the more emissions-free energy we can harness to power next-generation cement kilns and steel mills, the more carbon we'll be able to remove from the lifecycles of our most indispensable products. The overall goal should be to generate energy sustainably—and then transform commodity markets by embedding that energy into useful things.

Figure 2.7
GHG Emissions from Steel Production

Most steel is made using highly polluting ore-blast furnace technology (primary production). A smaller amount is made using electric arc furnace systems (secondary production). But as steel production has grown around the world, the total greenhouse gas output of both processes has continued to rise.

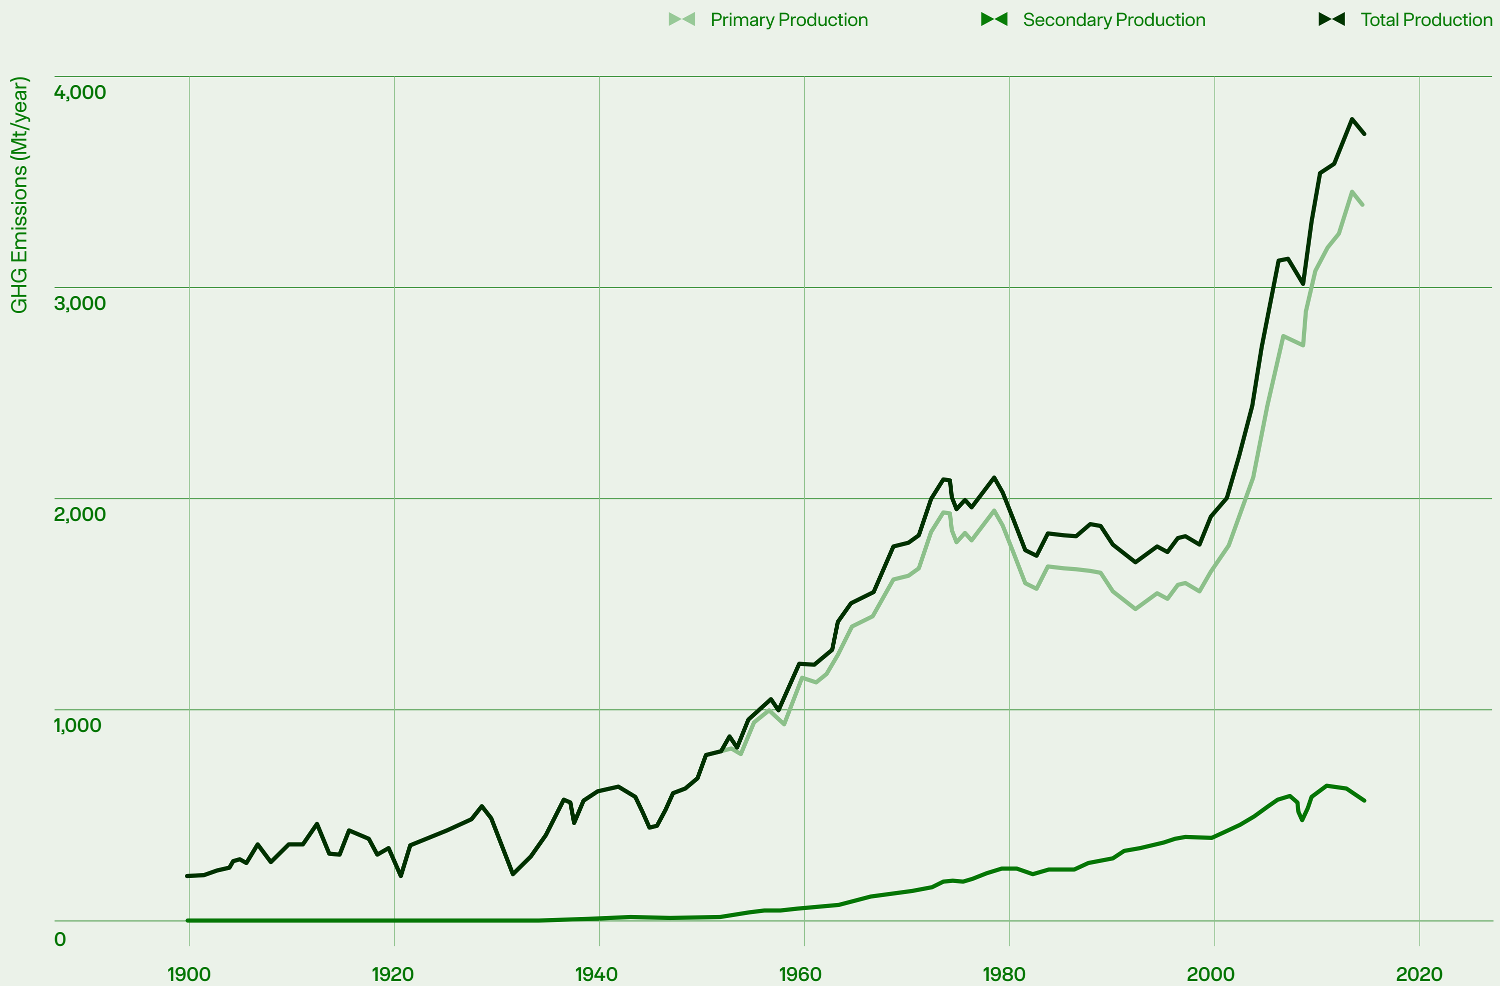

3.0

Computational Biology and Health

A complex organism such as a human is made of components that self-organize at different scales. Oxygen, carbon, hydrogen, and nitrogen, which account for 96 percent of our mass, combine with a few other elements into organic molecules like amino acids, which form proteins, which form cells, which form blood and tissues, which form our bodies.

The investment strategy at DCVC Bio, a life sciences specialist fund launched in 2018, spans the same range of size scales—as do the three contentions shared by the fund's managing partners, Dr. John Hamer, Dr. Kiersten Stead, Matt Ocko, and Zachary Bogue.

We look for entrepreneurs using new computational techniques to make predictions at the smallest size scales—for example, whether candidate small-molecule and protein drugs will be safe and effective. We're backing companies working at the microbial level to map the bacteria living in the human gut and help plants produce their own fertilizers and pesticides. And one startup backed by DCVC Bio is working at the cellular level to reprogram the human immune system to attack and clear solid tumors.

Add it up, and you get big advances in the way we do things at the human scale—whether that means caring for cancer patients or growing more plentiful crops and materials, all with fewer inputs and less pollution.

▸ We will use AI to design drugs we know in advance to have low toxicity, speeding up clinical trials.

Featured Companies

Creyon Bio [DCVC]
AbCellera [DCVC]
Totus Medicines [DCVC]

Computational Biology and Health: Predictive Drug Development

Experienced drug hunters are learning how to apply machine learning and proprietary experimental assays to the design of therapeutic molecules that have predictably low toxicity, potentially expanding the pipeline of treatments for both common and rare diseases.

The code of life is written in nucleotides. Strung together, they make up both DNA, which stores our genes, and RNA, which carries genetic instructions into the cell for translation into proteins. As soon as we learned how to read this code, we began to look for ways to rewrite it and control its transmission. And when drug developers began to figure out how to synthesize and deliver short strings of nucleotides called oligonucleotides, we entered the era of "information drugs" engineered to overcome specific flaws in the body's genetic program.

Antisense oligonucleotides (ASOs) and small interfering RNA (siRNA), for example, can prevent defective genes from producing disease-causing proteins, by destroying or repairing the intermediary messenger RNA. To date, only a handful of antisense treatments and other oligonucleotide-based medicines (OBMs), developed by industry pioneers such as Ionis, Alnylam, Moderna, and BioMarin, have won regulatory approval. But the possibilities of this approach are limitless. Once we understand how a mistake in our DNA or RNA causes a disease, we can, in principle, design an OBM to fix it.

There's just one big problem. Inside the body, many potential drug molecules do more harm than good. That's why traditional pharma companies must sift through thousands of drug candidates, and spend months or years on toxicity studies in animals, to find just a few molecules that don't have dangerous side effects. And it's why the FDA approves only 50 or so novel drugs each year, while millions of people with rare genetic conditions go untreated.

What's needed is a new generation of technologies for engineering compounds that drug makers *know in advance* will be safe and effective, eliminating the old trial-and-error approach to drug screening and vastly speeding up safety studies and clinical trials. DCVC Bio has invested in a collection of startups exploring new computational approaches to this problem, for both OBMs and traditional small-molecule drugs.

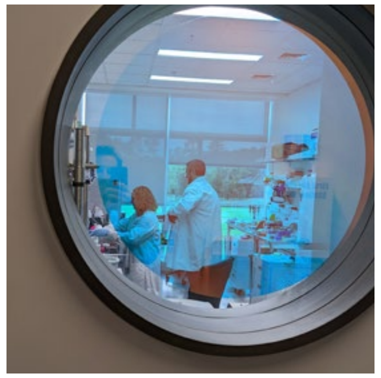

▸ **Creyon Bio**
Using bespoke datasets, Creyon's machine learning models uncover the chemical design rules that allow the company to create optimized oligonucleotide-based medicines (OBMs).

"There's a lot of toxicity associated with RNA medicines, and right now the entire industry just keeps making different ones to see which one is going to work," says Hamer, a PhD molecular biologist who formerly ran Monsanto's venture investing. "It's sort of a wreck-and-check approach—as opposed to being able to use, at the outset, a design process that gets rid of the things that are going to be toxic. We've got companies that really at the bleeding edge of trying to do this."

One of those is Creyon Bio, which takes a combined *in vivo, in vitro,* and *in silico* approach to reinventing drug design and screening. The company has built a library of survey compounds representative of the kinds of nucleotide sequences and added chemical units found in OBMs. It studies how those compounds affect the livers, kidneys, cardiovascular systems, and nervous systems of mammals. At the same time, test-tube assays reveal whether the survey compounds are effective in patient cells. The result is a proprietary dataset so large that the company can use it to train machine-learning algorithms to predict whether novel OBMs will have toxic or immune-stimulating side effects.

"We're trying to solve the oligonucleotide design problem, which is that the space of possibilities is so huge, trial and error completely fails," says Swagatam Mukhopadhyay, Creyon's cofounder and chief scientific officer. For an antisense drug with just 16 nucleotides, he notes, the number of potential variations is in the billions. "Forget screening—even computationally, you couldn't explore that space efficiently. So we said, 'What is the minimal set of sequences we need to study to be maximally informative about the pharmacologic design space?' That generalization allows us to come up with molecules that are orders of magnitude better, in the sense that we have high confidence in their safety and efficacy from the get-go."

The consequence is that Creyon can predict both the safety and the effectiveness of a molecule before it's ever delivered to a patient. The approach should work not just for antisense drugs but for other classes of OBMs such as RNAi, siRNA, ADAR guides, gRNA, RNase H-engaging molecules, and aptamers. Creyon intends to use its technology to dramatically accelerate OBM drug development and leapfrog competing drug makers to market. That could make it far more practical to design medicines for the more than 7,000 disorders that are linked to rare genetic mutations and have no available treatment.

The key, however, is to begin with relevant and highly curated training data. “There’s a lot of misunderstanding in the industry that you can take a drug development pipeline, sprinkle some AI on top of it, and it’ll make everything magical and more predictive—and that’s just not true,” says Stead. “AIs *are* superhuman, but they’re also incredibly narrow, and they have to be purpose-built to solve a particular problem. In Creyon’s case, it’s going to allow them to develop drugs for populations that have never had a therapeutic journey.”

Creyon is far from the only startup in the DCVC Bio portfolio applying machine learning to drug discovery. AbCellera uses the technique to search antibody structure/function relationships to find novel drugs in record time. Totus Medicines is using high-throughput cell-based screening and machine learning to discover new covalent drugs, which can irreversibly bind to disease-causing proteins and silence them. Empirico, in partnership with AbCellera, is integrating genetic data and health records for millions of individuals to discover and validate drugs that target previously unknown relationships between genetic variation and disease risk.

“The applications of machine learning and artificial intelligence in the pharmaceutical sector are still very much in their infancy,” Hamer says. But with a rational, predictable way to design low-toxicity medicines, he says, drug makers could create treatments for the millions of Americans with rare, untreated diseases: “We call that the long tail of drug development, and no one’s touching it today.”

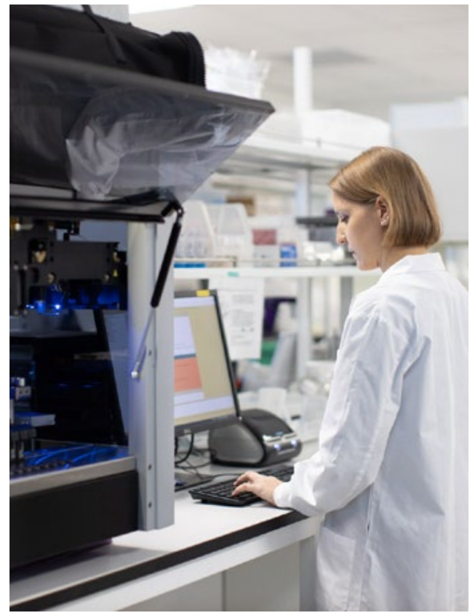

▸ AbCellera
Just months into the pandemic, AbCellera used its advanced machine vision and AI tools to identify an antibody against the SARS-CoV-2 virus.

▸ Totus Medicines
By screening billions of small molecules against thousands of cellular targets in a single experiment, the company finds effective new drugs thousands of times faster than traditional drug discovery and at a fraction of the cost.

“

AIs *are* superhuman, but they’re also incredibly narrow, and they have to be purpose-built to solve a particular problem.

Kiersten Stead
Managing Partner, DCVC Bio

Figure 3.1
Drug Development: R&D Spending and New Drug Approvals

For roughly a decade, from 2000 to 2010, U.S. drug industry R&D spending increased every year without any corresponding increase in the number of new drugs approved by the FDA. Since then, the pace of drug approvals has picked up again, but is still very low, at less than 50 new medicines per year.

Source: Congressional Budget Office

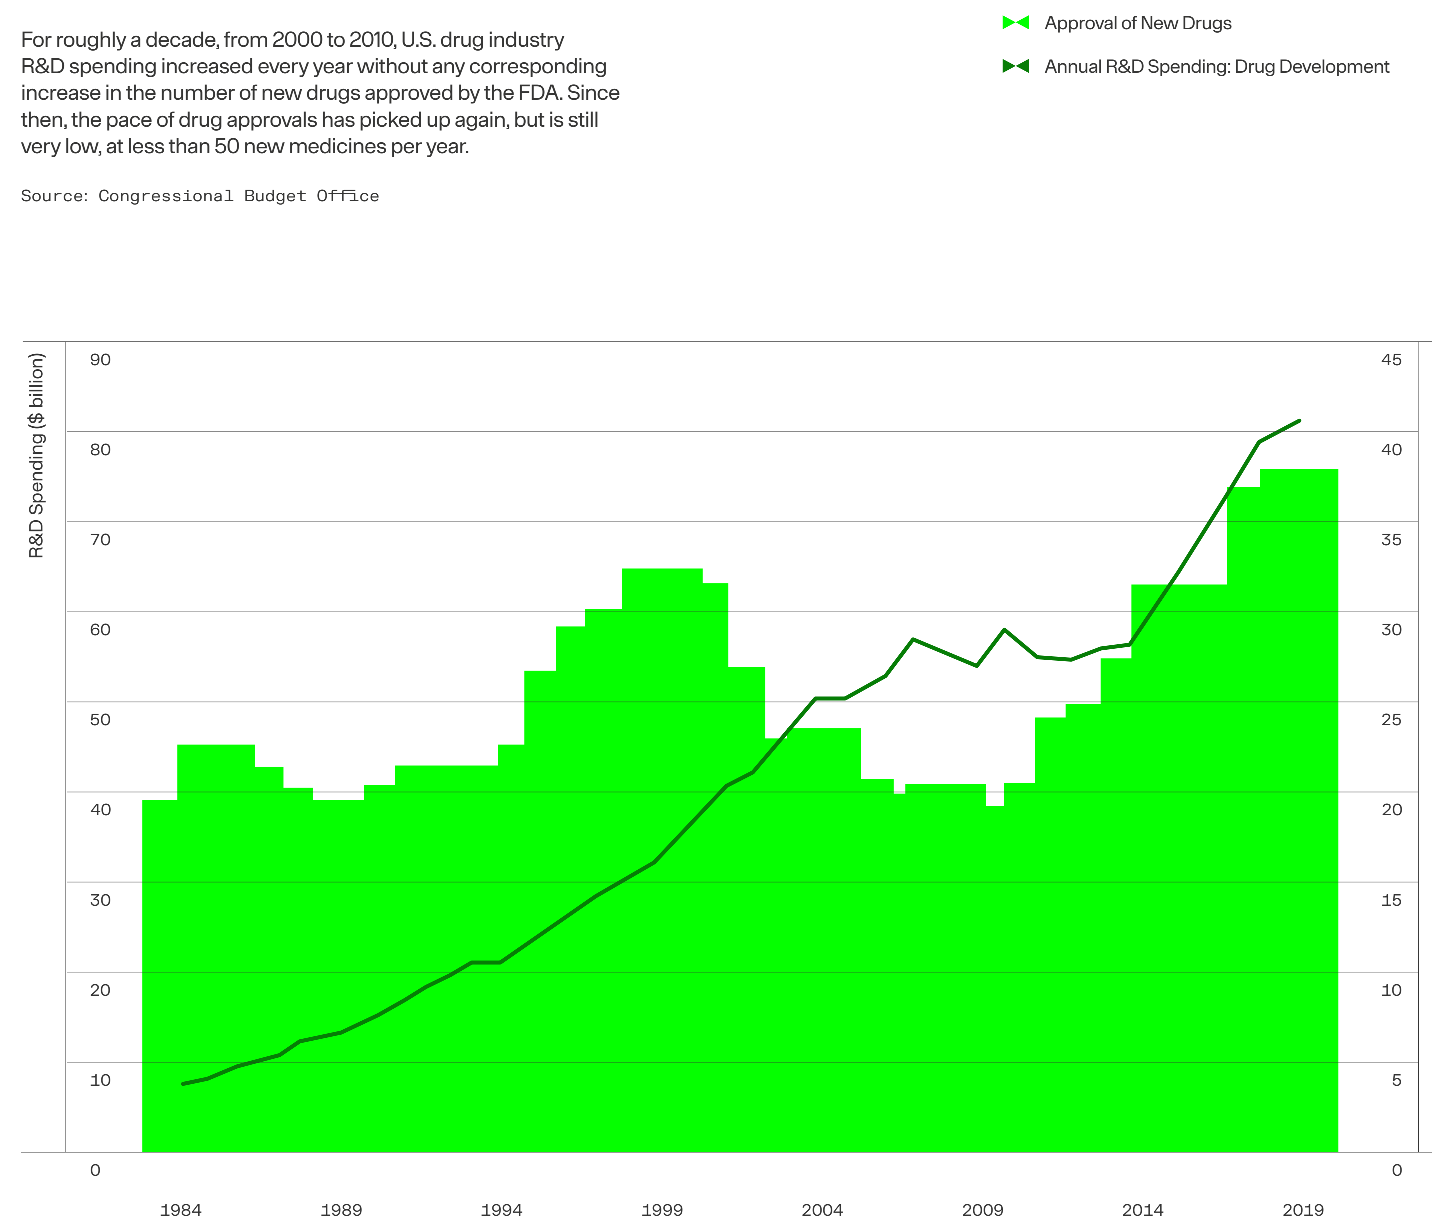

▸ We will turn our own cells into personalized drug factories, and learn to map and manage our microbiomes.

Featured Companies

Umoja Biopharma [DCVC]
Kanvas Biosciences [DCVC]

Computational Biology and Health: The Engineered Patient

We can now reprogram our own cells to manufacture drugs inside our bodies. In the early days of this technology, the reprogramming occurred exclusively *in vitro*, in the lab, but we're rapidly developing the ability to do it *in vivo*, in patients. We can put new genes into immune cells that turn them into tumor fighters, or map the microbiome as a step toward managing it more intelligently. These novel techniques could offer new hope to hundreds of millions of people with cancer, autoimmune conditions, and rare genetic diseases.

The coronavirus pandemic made history in many ways, and one of them was that billions of people received vaccines carrying a wholly new kind of cargo: synthetic messenger RNA. These threads of nucleotides carried instructions that showed our ribosomes, the protein-synthesis machines of our cells, how to express proteins resembling SARS-CoV-2's spiky coating.

In other words, scientists proved—on the largest scale imaginable—that our own cells can be repurposed as microbial protein factories, teaching the body how to defend against viral invaders.

But mRNA vaccines are only the beginning. Commercial oncology treatments that use genetic modification to convert our immune cells into cancer fighters are already available, and now biotech innovators are pursuing similar techniques that could permanently correct the genetic changes that cause many common and rare conditions. The nascent field of *in vivo* cell engineering is one of the opportunity areas where DCVC Bio is concentrating its resources and expertise.

"We're investing in treatments that are one-shot, treatments where *you* produce that drug for life," says Stead, who earned a PhD in molecular biology and genetics. "They're using you as the manufacturing vehicle."

Some of the original leads in cell engineering came from the field of cancer immunotherapy. More than two decades ago, researchers began designing artificial versions of T-cell receptors: the proteins on the surfaces of some white blood cells that recognize antigens and initiate an immune response. They hoped T cells equipped with these chimeric antigen receptors (CARs) would recognize, target, and destroy cancer cells.

The idea worked. There are now six FDA-approved CAR-T cell therapies for blood cancers such as lymphoblastic leukemia, B-cell lymphoma, and myeloma. They are remarkably effective, helping many patients achieve lasting remission.

But CAR-Ts are among the most expensive therapies in medicine, ranging from $350,000 to $1 million or more per patient, and have many complications. It can take several weeks to remove a cancer patient's own white blood cells, transfect the cells with the DNA encoding the required chimeric antigen receptor, and then grow enough of those modified cells for retransfusion (a process called expansion). Meanwhile, patients must also undergo lymphodepletion: chemotherapy that kills most of their remaining white blood cells, leaving them vulnerable to infection. On top of all that, CAR-T transfusions can cause life-threatening side effects such as cytokine storms.

It would be faster, safer, and cheaper to teach a cancer patient's T cells how to make tumor-targeting receptors without ever removing the cells from the body. And that's exactly the approach taken by Umoja Biopharma, one of DCVC Bio's portfolio companies.

Umoja encodes the instructions for a receptor targeting cancer cells into RNA. It puts that RNA inside a protein capsid shell and an outer lipid envelope—a vector system inspired by the life cycle of the human immunodeficiency virus (HIV). Once the vector enters a white blood cell, the RNA is reverse-transcribed into DNA and integrated into the cell's genome, where it directs expression of the chimeric receptor and equips the cell to attack tumors. All this happens without extraction, expansion, transfusion, or lymphodepletion. The process can even be turned on and off at will: Umoja tailors the new RNA so that the

▸ Umoja Biopharma
By reprogramming patients' immune cells inside their own bodies, rather than extracting and re-transfusing those cells as in CAR-T therapy, Umoja can attack cancers faster and at lower cost.

“We’re investing in treatments that are one-shot, treatments where *you* produce that drug for life ... They’re using you as the manufacturing vehicle.”

Kiersten Stead
Managing Partner, DCVC Bio

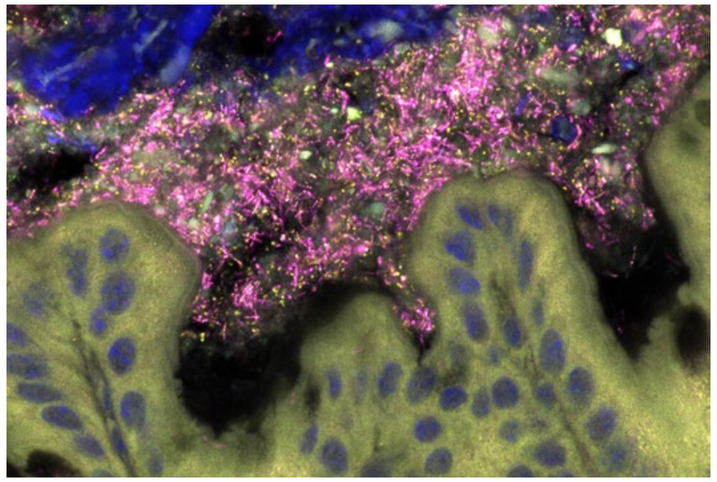

▸ **Kanvas Biosciences**
The company’s advanced tagging and classification technology will help us to map and control the populations of microbes in our bodies.

receptor-making genes are activated only when a safe, approved small-molecule drug related to rapamycin is present.

The company’s scientists believe this approach will lower the overall risks and costs of cancer treatment, and provide a built-in hedge against recurrence. “Once those cells are created, they’re engineered inside you forever,” Stead explains. “This old paradigm of going into remission, coming out of remission, and having to increase the aggression of your cancer treatments is going to be something that we can get around, because we will just re-induce that cell system through the action of the small molecule.” Similar technologies will likely work against other immune-mediated diseases, including lifelong autoimmune conditions like systemic lupus erythematosus.

DCVC-backed Kanvas Biosciences is exploring ways to control a different population of cells. The human gut and other epithelial regions, such as the vaginal tract, are home to thousands of species of bacteria, viruses, fungi, and other microorganisms. (In fact, the walking colonies we think of as our bodies contain more non-human cells than human cells.) Many of the bacteria perform vital services, such as breaking down indigestible materials, excreting essential nutrients, or protecting us against infections. But when their populations are out of balance, health crises can result. Kanvas is using RNA mapping of both microbial and host cells to figure out not just which microbes are present in our bodies, but exactly where they are and what effects they’re having.

In one experiment, for example, Kanvas was able to use its multiplexed fluorophore tagging system for RNA—combined with spectral imaging and a form of machine learning—to show that a new cocktail of beneficial bacteria is able to protect against an overpopulation of *Gardnerella vaginalis*, a type of bacteria that can cause vaginosis. Kanvas CEO and cofounder Matthew Cheng shared a dramatic micrograph showing how Kanvas’s cocktail preferentially adhered to vaginal epithelial cells, virtually eliminating *Gardnerella* from the sample.

“The microbiome is not only critical to human health, but should be considered a separate, distinct organ system within the human body,” Cheng says. “And because of our ability to look not only at what’s happening in the microbiome, but also in real time see what’s happening in the host, we believe we have a unique opportunity to design the next generation of microbiome therapeutics—to see how they engraft, look for safety, and look for efficacy, all in the same assay.”

The more researchers learn about how to alter gene expression in cells, how to get genetic payloads into the body, and how to work with (rather than against) the human immune system and microbiome, the sooner they will be able to intervene against a range of stubborn health problems. Already, companies are experimenting with cell therapies that create customized T-cell receptors designed to target the unique mutations in the tumors of cancer patients; separate out the therapeutic cells in stem-cell transplants from those likely to provoke rejection; and even help people with hair loss grow new follicle cells. “This is a trend that is in flight or at the beginning of being in flight, and I don’t think there’s a lot of visibility into it” among investors, says Stead. “But it’s something we think is going to be transformational.”

▸ On farms, we will learn to reduce the use of harmful synthetic fertilizers and pesticides by mimicking nature.

Featured Companies

Pivot Bio [DCVC]
BioPhero [DCVC]
MycoWorks [DCVC]

Computational Biology and Health: Redirecting Nature

To protect the environment and reduce harms to farm workers, we must vastly reduce our use of chemical fertilizers, pesticides, herbicides, and land, even as we find new ways to boost total agricultural production.

One answer is to collaborate more closely with nature itself, harnessing microbial products, insect pheromones, and other biological mechanisms to keep fields safe, productive, and pest-free.

The technologies that allowed us to expand farm production over the last century to feed a world population of eight billion people are not the same technologies that will allow us to feed an expected nine to 12 billion by the year 2100. One big obstacle is that today's farms depend on large volumes of synthetic nitrogen fertilizer.

Billions of people would be starving today—or would never have been born at all—if the German chemists Fritz Haber and Carl Bosch hadn't discovered a way to pull nitrogen out of the air in 1909. The tradeoff is that the Haber-Bosch process requires enormous amounts of natural gas as feedstock and fuel, generating between 1 and 2 percent of global CO_2 emissions. And much of the fixed nitrogen that we pour on our fields either cycles back into the air, forming polluting nitrogen oxide (NO) or nitrogen dioxide (NO_2), or runs off into rivers and lakes, creating algal blooms and huge oxygen-depleted dead zones.

Conveniently, some soil microbes have their own ability to fix nitrogen. Pivot Bio, which has been backed by DCVC since its seed stage in 2011, took on the challenge of finding bacteria that are especially good at generating natural ammonia (NH_3) fertilizer when added to seeds or to the rhizosphere, the envelope of soil around crop roots. Once it identified promising strains, the company also found genes in the bacteria that interrupt natural ammonia production in the presence of synthetic fertilizer, and turned them off. The result was a range of natural soil additives that work in concert with synthetic nitrogen, allowing farmers to reduce fertilizer use on corn, wheat, and sorghum crops by 30 percent with no drop in yields.

That's a major win for the environment and the climate. And in the future, the company aims to produce microbial additives that are even better at offsetting fertilizer use, in more types of crops. "Pivot is close to being the most transformational climate company I can think of," says DCVC managing partner Zachary Bogue.

A more recent DCVC Bio investment, BioPhero, is piggybacking on nature in a different way. Insects eat up to 20 percent of crop food mass before it even leaves the field. Historically, farmers have tried to reduce this toll by spraying plants with chemical pesticides. But these compounds are tightly regulated and are falling into disfavor, and in some cases the target species evolve resistance.

Starting in the 1970s, winemakers plagued by pests like *Cochylis* moths began experimenting with a sly alternative: inundating fields with lab-made versions of the pheromones emitted by female insects to attract

▸ **Pivot Bio** is replacing much of the synthetic nitrogen fertilizer that sustains today's crops with nitrogen fixed from the air by bacteria added to seeds or the soil around crop roots.

Figure 3.3 Global Fertilizer Prices

Alternative nitrogen sources could insulate farmers against wild fluctuations in the cost of synthetic nitrogen fertilizers, which use methane as a feedstock. The price of nitrogen fertilizer (in the form of diammonium phosphate) has tripled since 2020, as has the price of potassium-based fertilizers—largely thanks to the pandemic, the Ukraine conflict, and the resulting supply chain disruptions and natural gas shortages.

Source: Green Markets, Bloomberg; World Bank Commodity Markets

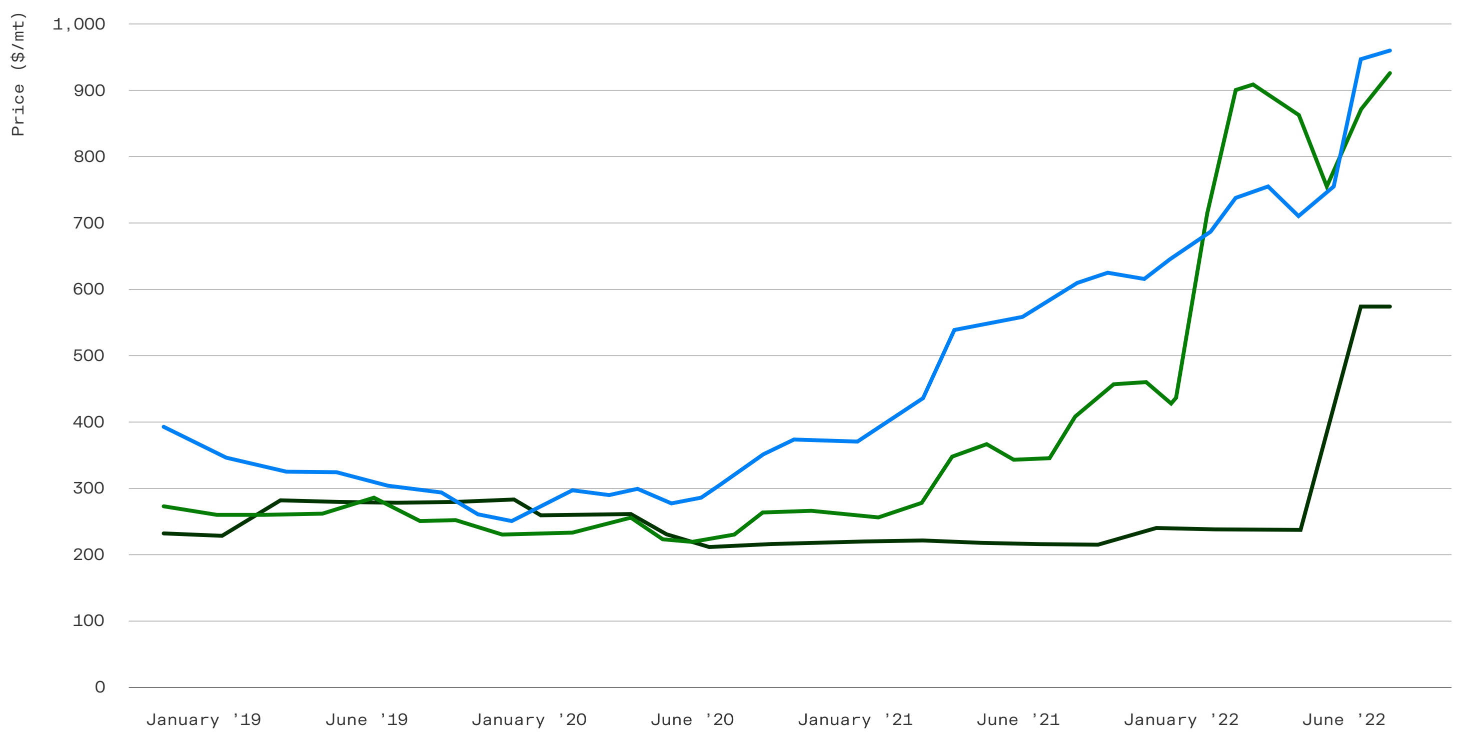

mates. The method worked, confusing males with overwhelming or false sensory input, with the result that moth populations plummeted within a single generation.

Pheromones have the advantage of being nontoxic for humans, and because they're critical to the insect life cycle, bugs are unlikely to evolve resistance to them. There's only one obstacle to scaling up this mating disruption. "To make an insect pheromone through traditional chemical synthesis is extremely expensive," explains Hamer. "So you can only use it on supervaluable crops like strawberries and orchid gardens. But where we really need it is in major crops."

That's what BioPhero is working on. The company's founding scientist, Irina Borodina, came up with a way to reprogram yeast cells to make pheromone components at industrial scale in bioreactors, similar to those used by biopharma manufacturers to make recombinant drugs. Today, the company is testing its synthetic pheromones as a weapon against the cotton bollworm *Helicoverpa armigera* and the fall armyworm *Spodoptera frugiperda*, which feasts on a range of crops such as corn, rice, peanuts, apples, and oranges. DCVC Bio led BioPhero's Series A funding round in 2021, and agricultural sciences company FMC acquired the company in 2022.

Even if they're made by yeast cells, BioPhero's pheromones will likely cost more per gram than chemical pesticides. But that's okay, because they're also much more potent. "If you can use 98 percent less volume, you can start looking at pretty interesting chemistries, because all of a sudden you don't need it to cost $2 a gram," says Stead. "It can be $10 a gram, and you can use it in a very prescribed way."

Entrepreneurs are also finding ways to engineer existing organisms to replace some of the products of livestock agriculture. DCVC Bio–backed MycoWorks is using mycelium, the tight networks of rootlike structures formed by fungi, to make an improved alternative to leather.

Naturally growing mycelium forms a foamlike material that can be compressed to make what's been called mushroom leather. But in this form, mycelium isn't as strong or durable as natural cowhide. MycoWorks has

engineered fungus growth to create intertwined fibers as they grow, and it can fine-tune features such as the material's thickness, softness, and flexibility.

In tanned form, fine mycelium from MycoWorks' pilot plant in Emeryville, California, is already used by Hermès for a new version of its Victoria bag, which is traditionally made from calfskin. Now MycoWorks is building a full-scale production plant in South Carolina where it hopes to supply leather alternatives for applications ranging from hats to car interiors. Other benefits of fine mycelium: it biodegrades at the end of its useful life, and producing it generates just 8 percent as much carbon per square meter as leather from cattle.

A century ago, we could not have imagined the depth of the environmental problems nitrogen fixation and other solutions for then-acute problems would create. But neither could we have imagined the breadth of the improved solutions we are inventing and implementing today. At DCVC, we believe in the wizardry of our scientists and entrepreneurs. Through a cycle of innovation and adaptation, we can preserve the natural world while making life better for everyone.

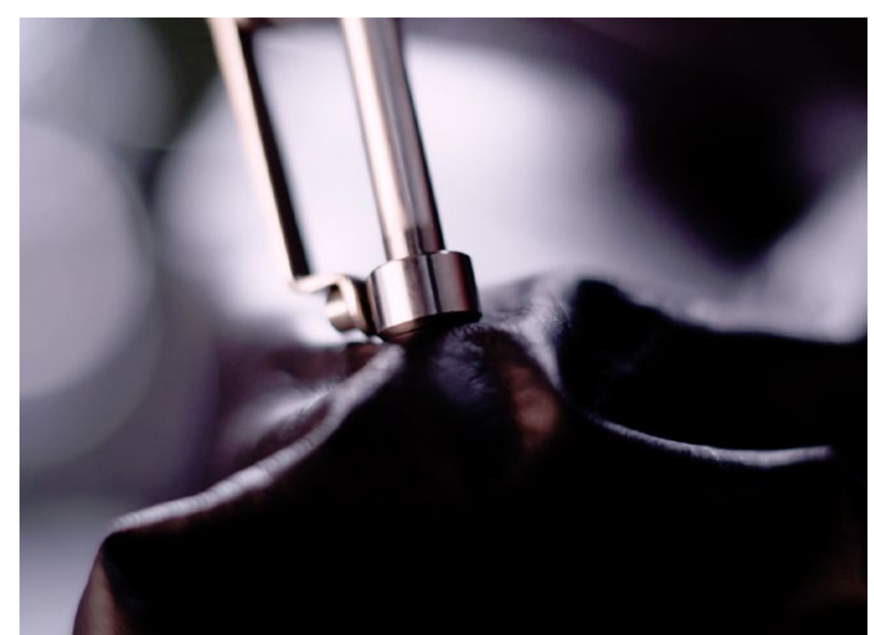

▸ **MycoWorks** engineers mycelium cells to create densely intertwined threads for high strength. Designers can specify properties not achievable with other synthetic alternatives to leather, or even animal leather. The company's Reishi™ brand material is already being used in hats, handbags, and other items.

Cellular Agriculture

Billions of dollars in capital is flowing into startups working on ways to grow meat in commercial quantities by cultivating animal cells in the lab. These companies have worthy goals: they aim to reduce the suffering of animals and blunt livestock agriculture's contribution to global warming. But for the foreseeable future, we believe, *in vitro* meat will remain in the research stages.

↳ The benefits of replacing animal-grown meat, especially beef, with other forms of protein are undeniably attractive. The typical cow or steer emits up to 500 liters of methane into the atmosphere every day, while bacterial breakdown of their manure adds even more of the potent greenhouse gas. Eliminating livestock production (assuming the existence of a pollution-free alternative protein source) would immediately cut global anthropogenic emissions of methane by 25–40 percent. And that's to say nothing of the CO_2 emissions that would be prevented if we stopped burning forests to clear land for ranching and cattle-feed production.

Unfortunately, the math behind growing and selling cultured meat as an alternative to meat from animals doesn't pencil out—not even remotely. "We're in venture, so we're never cynical enough to say that the scale and the price of a new technology doesn't change over time, because it absolutely does," says Stead. But in cellular agriculture, she says, there are countervailing forces that resist any economies of scale.

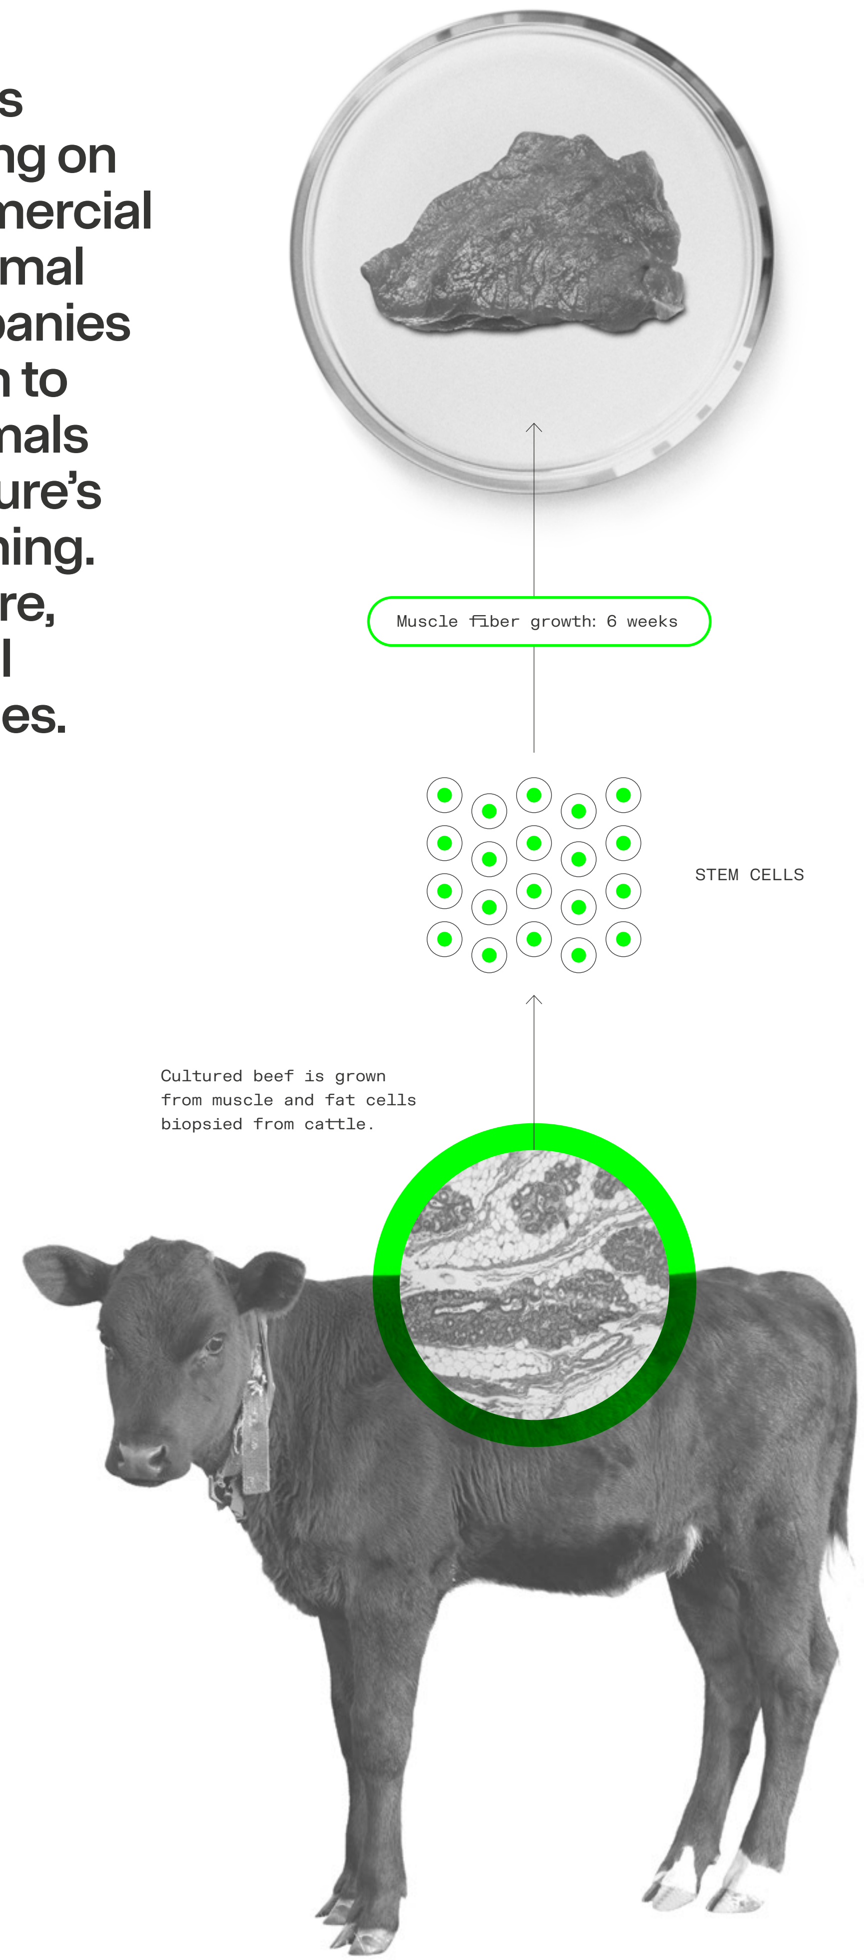

Figure 3.4.1

Eventually, cultivating meat in the laboratory could be more efficient and sustainable than raising cattle. But the process isn't as simple as growing a few cells in a Petri dish. The fleets of bioreactors needed to grow cultured meat at the scale required to make a real dent in today's food system would be ludicrously costly.

For example:

↓

Most companies aiming to make cultured meat on an industrial scale grow animal cells in bioreactors similar to those used to make vaccines and recombinant drugs. The Good Food Institute, a nonprofit representing the alternative-protein industry, says building a bioreactor facility large enough to produce 10,000 metric tons of cultured meat per year would cost $450 million. But annual meat production in the U.S. alone amounts to 12 million tons of beef, 12.5 million tons of pork, and 21.2 million tons of poultry, according to the USDA. To replace all of it, the cellular agriculture industry would need to build 4,500 factories, at a total cost exceeding $2 trillion. (That doesn't even count the cost of clean-room-style operations, or of producing the amino-acid-rich growth medium animal cells need to multiply.) And even if you could produce enough cultured meat, transporting it to end consumers safely would require a refrigeration chain that doesn't exist outside the developed world.

↓

Heating and cooling bioreactors requires large amounts of electricity. The consulting firm CE Delft, in an analysis commissioned by the Good Food Institute, found that the carbon footprint of a given unit of cultured meat actually exceeds that of chicken or pork, unless bioreactor facilities are powered using 100 percent wind, solar, or geothermal energy, which obviously limits the locations where they could be built. (Whatever the energy source, however, cultured meat has a significantly smaller footprint than beef.)

↓

Animals, especially chickens, are a cheap, abundant, efficient source of complete protein (the kind that contains the nine amino acids our bodies can't make on their own). Modern breeds of chickens convert 40 to 60 percent of their diet into protein that humans can consume—a much higher rate, at the moment, than cultured meat. Most people, moreover, are happily carnivorous; only in India do vegetarians and vegans exceed 20 percent of the population. It would likely take a global shortage of animal protein or massive price hikes to drive significant consumer demand for cultured meat, and no such crisis appears to be looming.

Meanwhile, there are ways to reduce the emissions profiles of conventional poultry, beef, and pork production that are far more economical than scaling up cellular agriculture. For instance, livestock can be fed deforestation-free soy and other more sustainable crops. Adding natural chemicals to cattle feed can reduce methane production in the guts of cows and steer by up to 90 percent (an innovation being pursued by DCVC/DCVC Bio portfolio company CH4 Global).

To Stead, these factors mean that cellular agriculture simply doesn't make sense as an investment opportunity on the time scales required by venture capital firms. "There is a strong desire to do the perceived good thing for the environment," she says. "But you need to weigh the realities of cellular agriculture on any meaningful scale against the benefits and efficiencies of animal agriculture, and the need for a global diet that includes the essential amino acids."

Figure 3.4.2 **Carbon Profiles of Meat Types** — CO_2 CH_4

The production of cultured meat is not substantially less carbon-intensive than conventional beef cattle production, according to a study led by University of Oxford researcher John Lynch. (The carbon footprints described in the table as a, b, c, and d differ mainly as a function of the design and operation of the relevant bioreactor and the specific feedstocks used to nourish the cultured cells.)

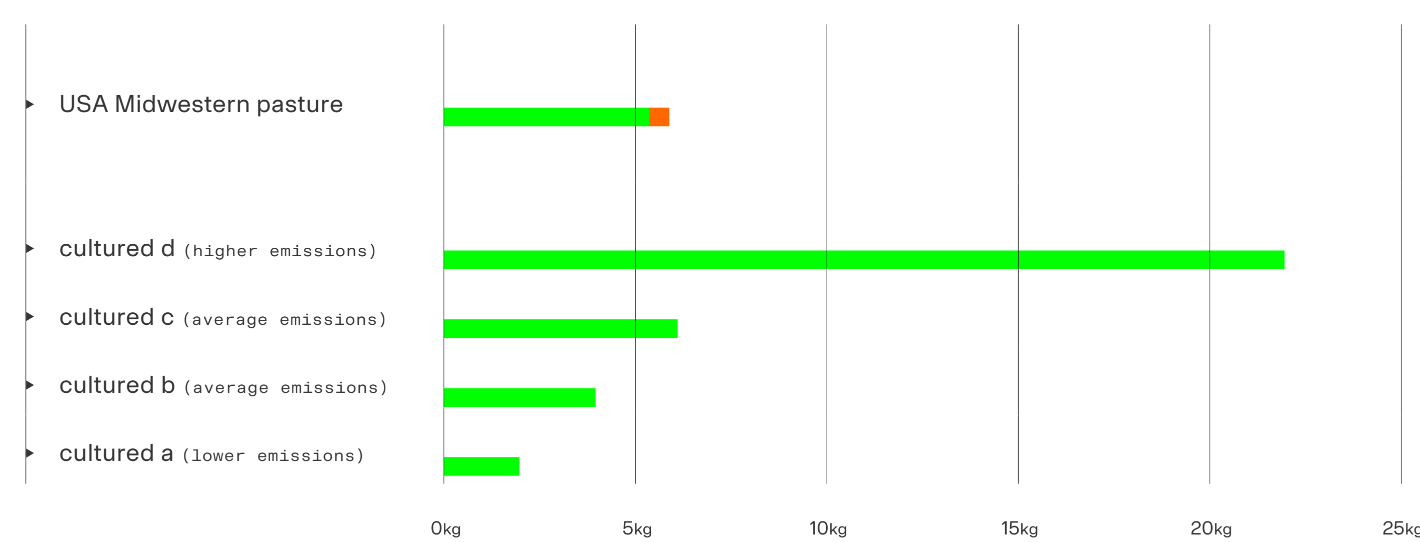

Source: Lynch, J. & Pierrehumbert, R. (2019). "Climate Impacts of Cultured Meat and Beef Cattle." Frontiers in Sustainable Food Systems, 3.

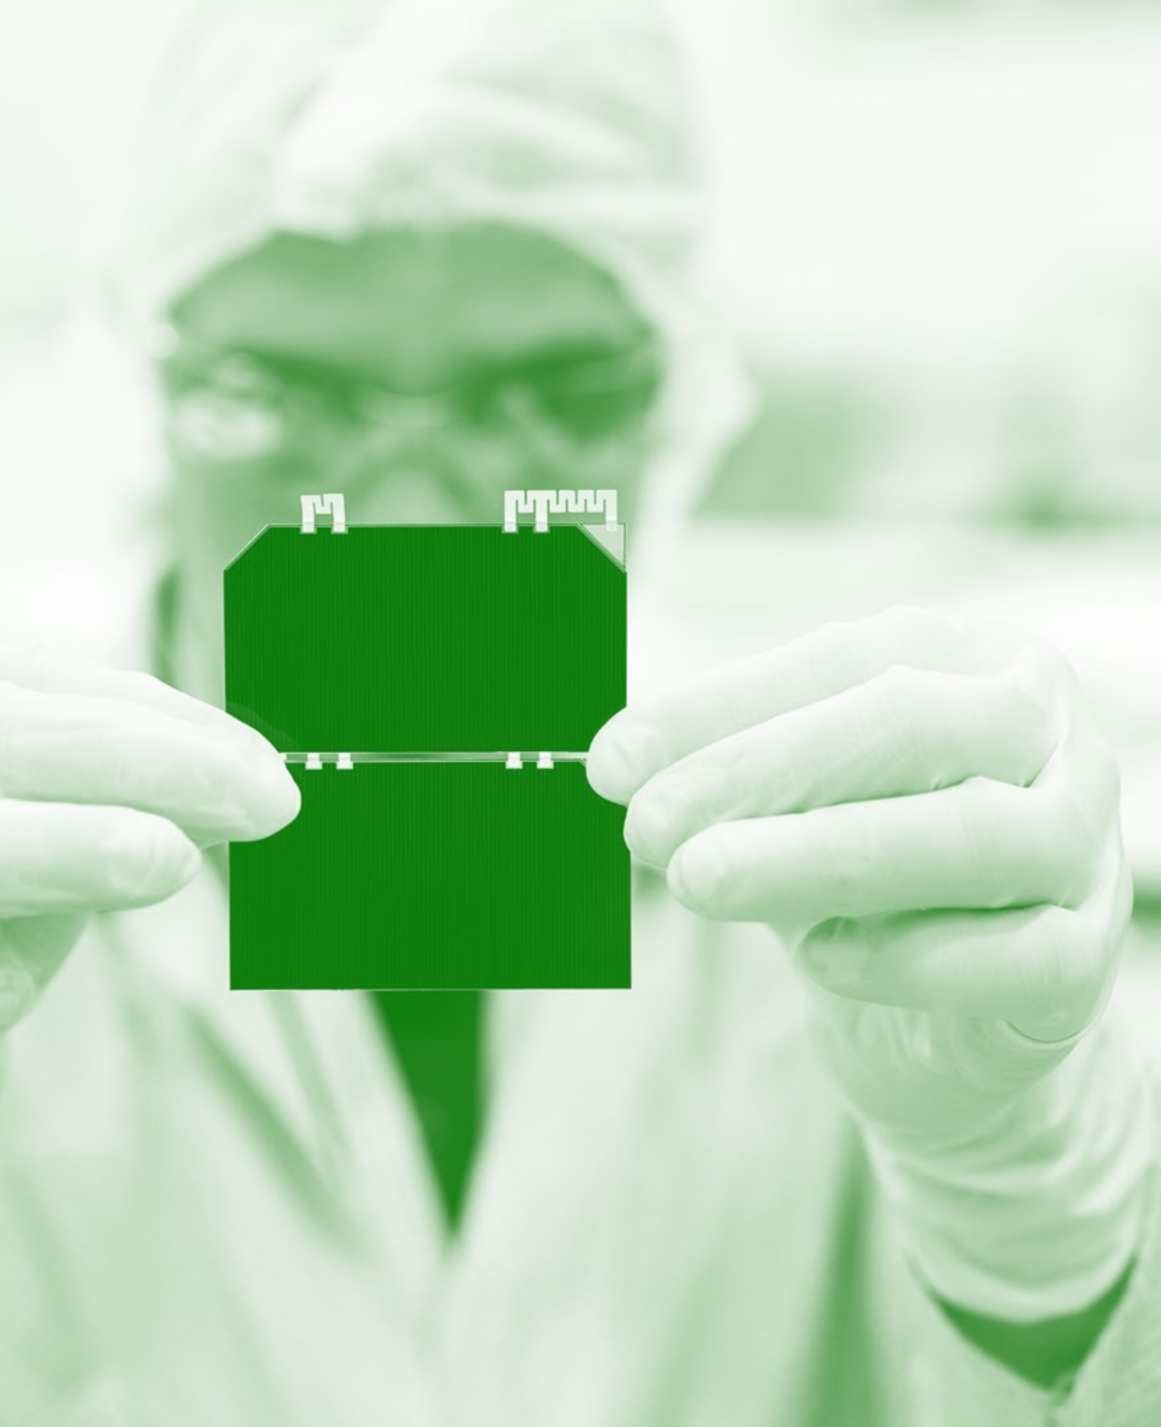

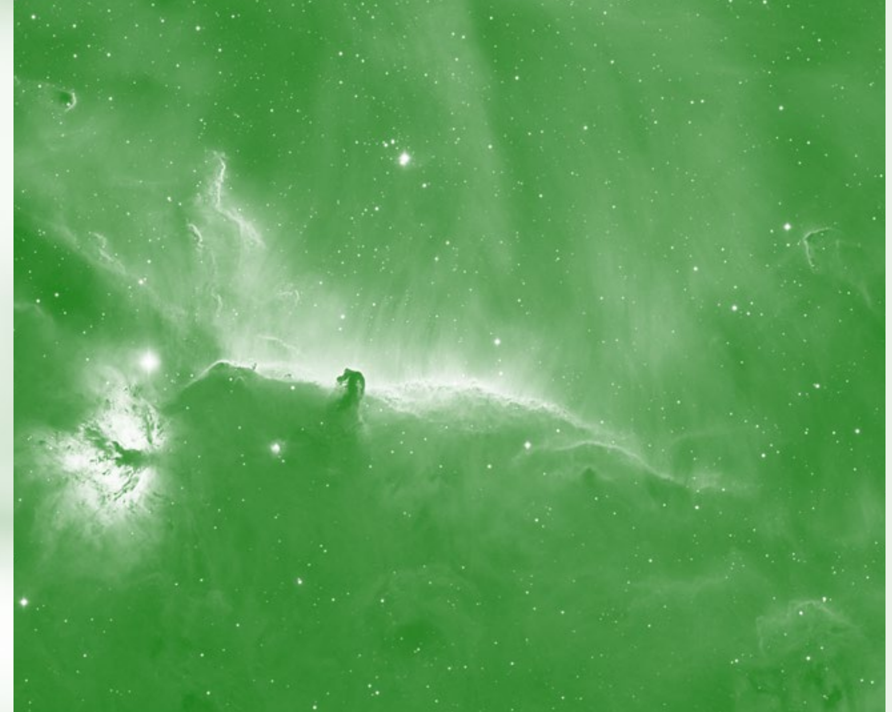

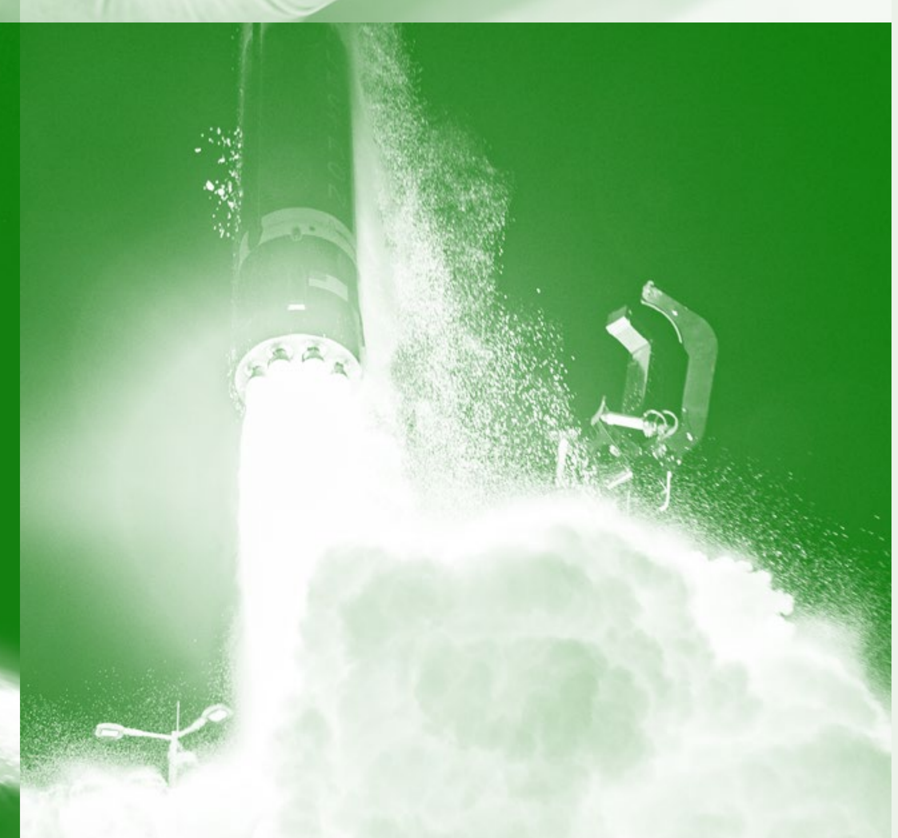

Space

The time has come for entrepreneurs to remake the space business. DCVC invests in companies that are putting satellite data to broader use, opening up access to low earth orbit, and making space more relevant to our everyday lives. “We’re at the point where smart engineers, working outside NASA or the defense industry, can help open up space to applications with far broader benefits than ever before,” says DCVC managing partner Matt Ocko.

The commercial space business is a lot older than most people realize: 2022 marked the 60th anniversary of the launch of Telstar I, the world's first commercial communications satellite. But the conditions for rapid startup innovation in space have only emerged in the last decade or so:

- With governments receding as the main customers for space hardware, aerospace companies moved beyond a cost-plus pricing regime—which had incentivized regular budget and schedule overruns—and began to accept more risk.
- Launch providers such as SpaceX figured out how to make rockets from cheaper components and bundle multiple satellites on a single rocket, dropping the cost-per-kilogram of getting to space drastically.
- Ever smaller electronics made it possible to design smaller, cheaper, and more capable satellites for communications, Earth observation, and other missions.

"We've evolved beyond the industrial age of space, when the focus was on bending hardware to achieve certain technical performance," says Matt O'Connell, a DCVC operating partner who previously led satellite companies GeoEye and OneWeb. "Now the focus is on bending wavelengths, bending pixels, and bending data."

The space companies we've backed at DCVC are working to further democratize access to space, gather more types of data, and make the next generation of space technology more resilient. TrustPoint, for example, is developing a fleet of microsatellites designed to make it faster and easier for devices on the ground to locate and lock onto a GPS signal, while also preventing impersonation by false signals. Rocket Lab has built a relatively small launch vehicle—the 59-foot-tall Electron (one-quarter the size of the SpaceX Falcon 9)—that has already delivered more than 150 small satellites to orbit. Its design incorporates a 3-D printed engine, reusable first stage, and unique kick stage. Planet and Capella Space both own constellations of satellites in low earth orbit that provide government and commercial customers with high-resolution digital images of the surface, aiding with everything from water supply monitoring to intelligence gathering to disaster recovery. Planet's images come in the form of conventional photographs in visible and hyperspectral wavelengths, while Capella uses synthetic-aperture radar that penetrates clouds, rain, smoke, and the dark of night. Both companies also operate cloud-based platforms for storing and managing terabytes of image data.

If there's a theme that unites these companies, it's that they're all going beyond an exclusive focus on space hardware to create services that improve conditions on the ground. "We see space as a gateway to all the other things we need to get done," Ocko says. "Monitoring greenhouse gas emissions, deforestation, and water use. Protecting property and national security. Understanding and facilitating the global movement of goods. It's all interconnected, and a world with a stronger space infrastructure is a happier and safer world."

Figure 3.5
Licensed Commercial Space Launches

The count of licensed commercial launches has increased dramatically since 2017, when SpaceX launched 18 Falcon 9 rockets into space.

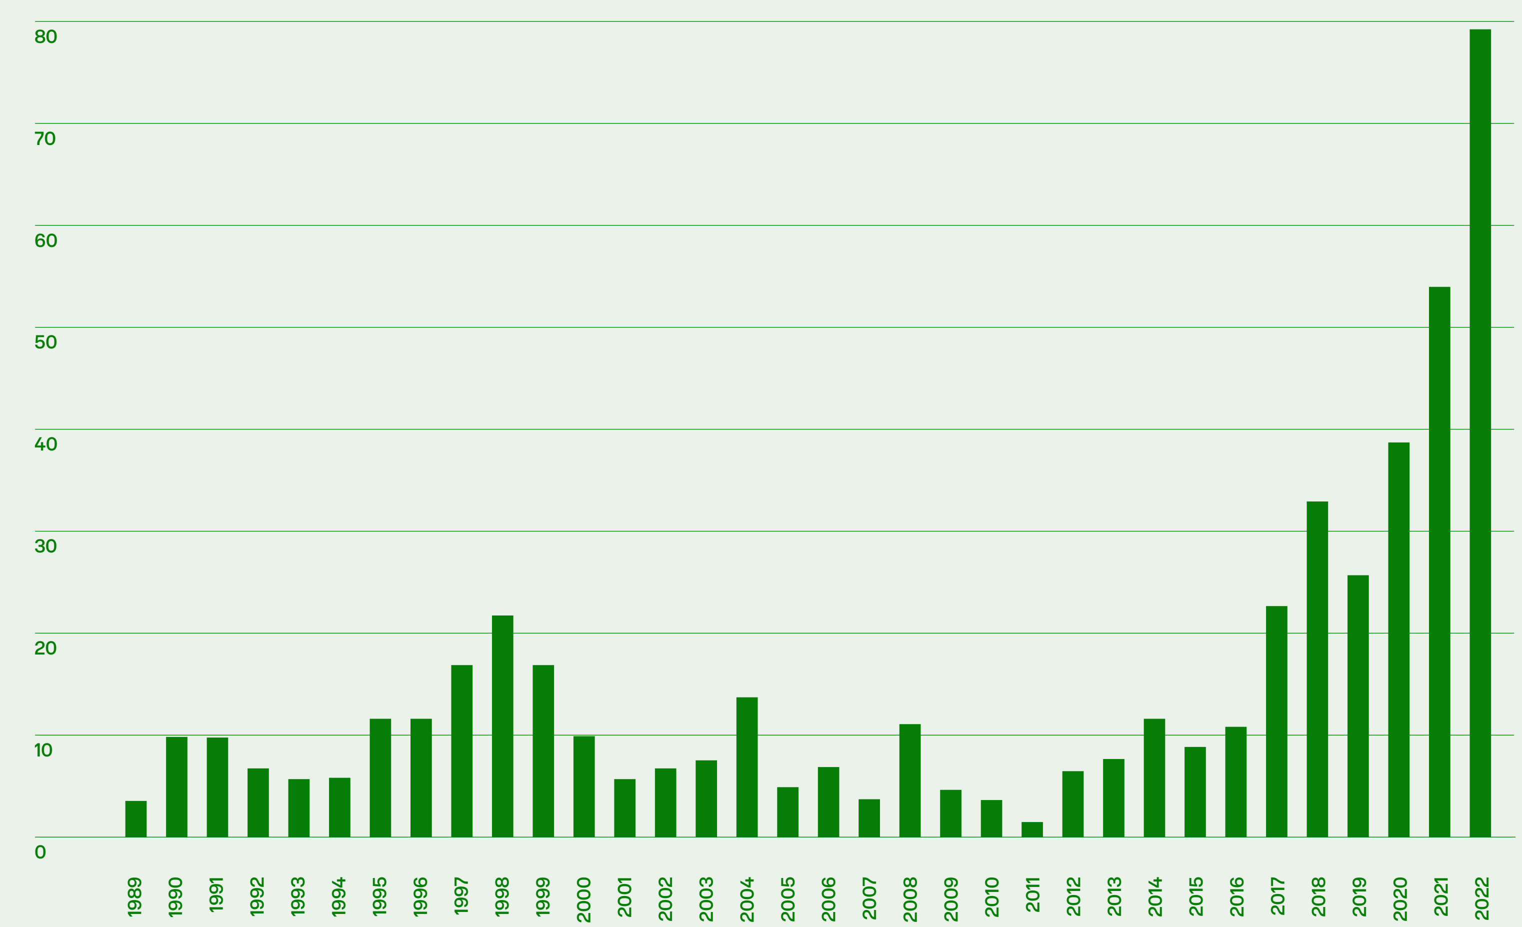

Source: Federal Aviation Administration https://www.faa.gov/data_research/commercial_space_data

Cryptography for the Real World

The internet was made safe by three critical waves of security innovation. The first came in the early 1990s with the invention of firewalls that protected private networks by restricting the movement of data in or out of the networks. Soon after, companies like Netscape and Verisign introduced Secure Sockets Layer certificates (succeeded by today's Transport Layer Security certificates), which used public-key encryption to authenticate web servers and other devices; this second wave ended up powering the commercial internet, from Amazon to healthcare portals to e-government. The third wave, being implemented today, is an upgrade to the infrastructure of trust—using blockchain and other techniques to ensure data integrity without compromising privacy.

Of course, it's no wonder that many consumers and business users are skeptical about the cryptocurrency world. When FTX and its native currency, FTT, collapsed in late 2022 amid revelations that its founders had allegedly used customer deposits to cover losses at FTX's sister trading firm, Alameda Research, it underscored a stubborn paradox. The blockchain itself, in which each transaction is verifiable by code, without the need for a central arbiter, is a powerful innovation that makes it easier not just to move money around, but also to write self-enforcing peer-to-peer contracts or employment agreements, trade carbon credits, fundraise for startups or nonprofits, or share authenticatable scientific information without extractive middlemen. But precisely because cryptocurrency markets are decentralized and largely unregulated, unscrupulous operators can abuse the trust of their investors or business partners. That is why DCVC is focused on innovation in *cryptography*, not in *cryptocurrency*.

The great promise in the idea of the blockchain was not that it would open up new avenues for financial innovation and decentralization—though it has done that, for better and worse. Rather, DCVC sees modern cryptography as the enabling technology for a new era of digital communications, shared scientific innovation, and commerce, where trust and verifiability are built in. Unregulated crypto markets are irrelevant to the real work of exploiting the capacity of the blockchain to build trust, being pioneered by companies using fundamental cryptographic principles. Companies in our portfolio are using those principles and related ideas from algorithmic finance to increase transparency and efficiency and reduce costs in critical areas of science and commerce, all in the service of creating equitable economic opportunities and new business models.

StarkWare, backed by DCVC since its 2018 Series A funding round, is an example. The company's cofounder and president, Eli Ben-Sasson, co-invented zero-knowledge scalable transparent arguments of knowledge, or ZK-STARK, a method for verifying the integrity of confidential, encrypted data without slowing blockchain-based markets to a crawl.

A zero-knowledge proof is a way to demonstrate the validity of a statement without revealing anything about the statement itself, and the idea has applications in fields as critical as nuclear disarmament, where it can be used to tell real warheads from spoofs without opening the devices to tampering or snooping. StarkWare is using it to overcome a limitation in first-generation cryptocurrencies such as Bitcoin. While Bitcoin funds are tied to pseudonymous addresses, not named individuals, every Bitcoin transaction ever completed is visible in its blockchain—meaning Bitcoin is anonymous but not private. Newer protocols such as Ethereum allow for the storage of hashed or encrypted data and transactions, but it takes a lot of computational power to verify that a hashed dataset is what its owner says it is. The ZK-STARK technique ensures privacy by shielding the inputs that go into a proof-of-integrity computation. And it solves the scalability challenge by vastly reducing the amount of computation needed to verify a batch of transactions.

StarkWare is building a set of tools that make it easy for developers to use ZK-STARK proofs on the Ethereum platform. One, called StarkNet, is a service that validates batches of transactions using ZK-STARK proofs before they're added to the Ethereum blockchain. Visa used it to test a way for citizens of multiple countries to make recurring payments automatically, essential for young businesses around the globe to make payments to their supply chain and to regulators. Another is StarkEx, an architecture for decentralized exchanges of value, including intellectual property (such as AI algorithms, data sets to train them, and the discoveries they may yield) and the funds to license and/or purchase it. StarkEx enables huge numbers of stakeholders to seamlessly rendezvous and exchange this value, in perfect trust, with zero fraud.

Collectively, StarkWare's capabilities deliver another level of trust, attestability, and privacy for valuable transactions on the Internet. "StarkWare is this decade's Verisign—a breakthrough in how large groups of people can trust each other and be secure in how they communicate things of value," says DCVC managing partner Matt Ocko.

Axelar is another DCVC portfolio company working to make blockchains more secure. There are already so many different blockchain platforms and projects that it's becoming a headache to move assets or contracts between them, leading developers to experiment with cross-chain bridges. Yet security flaws have exposed some existing bridges to hackers' attacks, which rules out their use for large corporations to exchange IP and other protected/attested information across their mutually distinct, corporate-specific blockchains. Axelar operates an overlay network that itself uses a decentralized design, plus a validation system called quadratic voting, to pass messages securely between different blockchains. For example, two large pharmaceutical companies can share information about a jointly developed drug securely, across their proprietary blockchain, allowing both to

validate—independently, cryptographically, and unalterably—which company contributed what to the final drug. "You can take a token that you have in one protocol and use it as collateral to basically perform something else on another blockchain," explains DCVC general partner Ali Tamaseb.

The work DCVC is supporting at StarkWare and Axelar goes "very deep, asking, 'How do we make the encryption more secure and scalable, how do we make it faster, how do we make the different layers and pipelines connect with each other better?'" Tamaseb says.

drugs will automatically adjust when the programs reach agreed milestones, such as proof of nontoxicity or a positive Phase I trial result. All the parties to the drug development process will be able to see and validate the data. And with science as king, promising drugs will no longer get sidelined purely due to poor management, poor marketing, or capricious capital markets. "If you really look at how we make drugs, it's not an emotional business," Paragas says. "It's brutally algorithmic."

And there are many other markets where transactions can be sped up, or opened up, through algorithmic thinking.

"You're using data to make the flywheel go faster, to build better models, to understand something better ... That's the deep tech lens.

Ali Tamaseb
General Partner, DCVC

A DCVC-backed company called DVLP Medicines is using blockchain technology to reinvent another essential function of our economy: the creation of new medicines. As more microbes evolve ways to evade our current antibiotics, there's a desperate need for a new generation of antimicrobial compounds. Yet such drugs don't usually command high prices, which gives Big Pharma few incentives to shepherd them through the risky, multiyear, multibillion-dollar process of getting from the lab bench to FDA approval.

But does the process *have* to be so costly? DVLP's cofounder and CEO, Jason Paragas, argues that traditional drug-development organizations are weighed down by too many layers of overhead. What's needed, he believes, is a way for investors to back specific drug candidates, rather than the companies around them. Inspired by the "payment rails" used by banks to track and verify digital money transactions, DVLP is creating a blockchain-based system that will allow drug developers to securitize specific antimicrobial molecules and other drugs with clinical promise. Because modern blockchains can incorporate smart contracts, the value of these securitized

"We believe in a future where many assets are tokenized," Tamaseb says. Kettle, for example, is updating the way the insurance industry uses statistical modeling to understand and price risk. The company's convolutional neural networks analyze weather data, vegetation, topography, building design, and other factors to predict where wildfires will start and how much damage they might cause. That allows Kettle to understand fire risk at a more granular level and offer insurers cheaper reinsurance. Tala does something similar in the world of small-business lending, using data from smartphones to assess creditworthiness and deliver automated loans to small business owners in Mexico, Kenya, the Philippines, and India.

In the end, DCVC's cryptographic companies are all developing smarter infrastructure, whether they are enhancing trust and safety, atomizing drug development, or finding better ways to model risk in insurance and lending. "You're using data to make the flywheel go faster, to build better models, to understand something better," says Tamaseb. "That's the deep tech lens."

“

Be a scientist, save the world.

Richard Smalley
Co-discoverer of buckminsterfullerene

The 2003 speech in which the Nobel Prize-winning chemist uttered this now-famous line was about the breakthroughs in physical science and engineering that will be needed to find new sources of clean energy. Which, to us, sounds like deep tech.

Many of the problems that will preoccupy us in this century—global warming, population growth, water shortages, fast-spreading pandemics, aging populations in rich nations and fast-growing populations in developing ones—are the unintended side effects of advances in technology since the first Industrial Revolution. We unshakably believe that the cure for the unwanted impacts of past innovations is not retreat, or worse, paralysis, but *more innovation*. It's the greatest game on Earth, and it never ends. If even half of the contentions DCVC's partners have shared here turn out to be correct, we'll all end up enjoying longer, more rewarding lives, on a planet where we have learned to sustainably and equitably manage the impacts of human activity.

Of course, it will take people with many kinds of skills—diagnostic, artistic, diplomatic, political—to make sure our species survives and succeeds in this game. But only better science and engineering can expand the game board itself.

If you are excited by the same opportunities we see—and if what we and our portfolio companies are working on seems compelling—we want to hear from you.

Thank you to the following DCVC partners

Zachary Bogue
Chris Boshuizen
Alan Cohen
Steve Crossan
John Hamer
James Hardiman
Earl Jones
Matt O'Connell
Matt Ocko
Jason Pontin
Rachel Slaybaugh
Kiersten Stead
Ali Tamaseb
Matt Trevithick

Editing Nate Nickerson
Research and Project Management Stephen Seymour
Research, Data, and Graphics Alex Gachanja
Writing Wade Roush
Copyediting Tim Farrington
Fact Checking Matt Mahoney
Photo Research Hailey Hiss
Design One Design Company, Chicago, IL

Typefaces FK Grotesk Mono, Kleukens-Egyptienne, Onsite
First Published July 2023

Deep Tech
Opportunities Report

2023

270 University Avenue
Palo Alto, CA 94301

500 2nd Street, Floor 2
San Francisco, CA 94107

dcvc.com