JOHN H. EVANS, CRAIG CALLENDER, NEAL K. DEVARAJ,
FARREN J. ISAACS, AND GREGORY E. KAEBNICK

Building Decision Points Into Research's Slipperiest Slopes

The controversy around a ban on "mirror life" should lead to a more nuanced public conversation about how to manage the benefits and risks of precursor biotechnologies.

About five years ago, the five of us formed a discussion group to investigate the ethics of working with the building blocks of *mirror life*—biomolecules that are essentially mirror images of their natural counterparts. We knew that less-novel research within the broader field of synthetic biology has been the subject of extensive social and ethical debates, so we wanted to be proactive about a technology we consider even more novel. Two of us (Devaraj and Isaacs) are synthetic biologists who focus on lipids, nucleic acids, and proteins. The rest are a philosopher of science (Callender), a social ethicist of biotechnology (Evans), and a bioethicist (Kaebnick). We met roughly quarterly to discuss what was happening in our labs, as well as the possible futures of such technology.

Last year, Devaraj and Isaacs were invited into a second group, a large consortium of dozens of synthetic biologists and research leaders preparing a statement on a topic our discussion group hadn't specifically considered: whether to make mirror bacteria. Out of this consortium came something remarkable. In December 2024, in a policy forum article in *Science* magazine, several of the world's most influential scientists called for a ban within their own research area.

Over the past half century, there have been calls for moratoriums in other research areas, such as human gene editing and human-animal chimera embryos (the latter ban was lifted in 2016). However, the call to prohibit the creation of mirror life, something that is currently impossible to do, was notable for being unequivocal. The authors of the article declared outright that "mirror bacteria and other mirror organisms should not be created." Their concern was that mirror cells could replicate unchecked by predators or immune surveillance, and so eventually outstrip natural counterparts to devastate ecosystems. The statement also called for the scrutiny of precursor technologies that could lead toward mirror life.

In the aftermath, our group has come to believe that the statement was strong but incomplete: It did not include what precursor technologies should be allowed. The five of us agree that nobody should create mirror bacteria, but we think concrete steps are needed to ensure a ban against doing so does not limit beneficial research. We realized the work of our discussion group could help provide a framework for how and whether to advance mirror life research. It is a strategy that titrates risks and benefits, aimed to allow aspects of mirror research to continue in a way that avoids the harms motivating a ban while still pursuing the benefits that motivate the science. Such a framework can also be useful in thinking about other potentially dangerous technologies.

A look in the mirror

The meaning of mirror life is tied to the concept of *chirality*, which comes from the Greek word for hand. Our left and right hands are basically mirror images of one another, yet a left-handed glove cannot fit a right hand. Chirality occurs

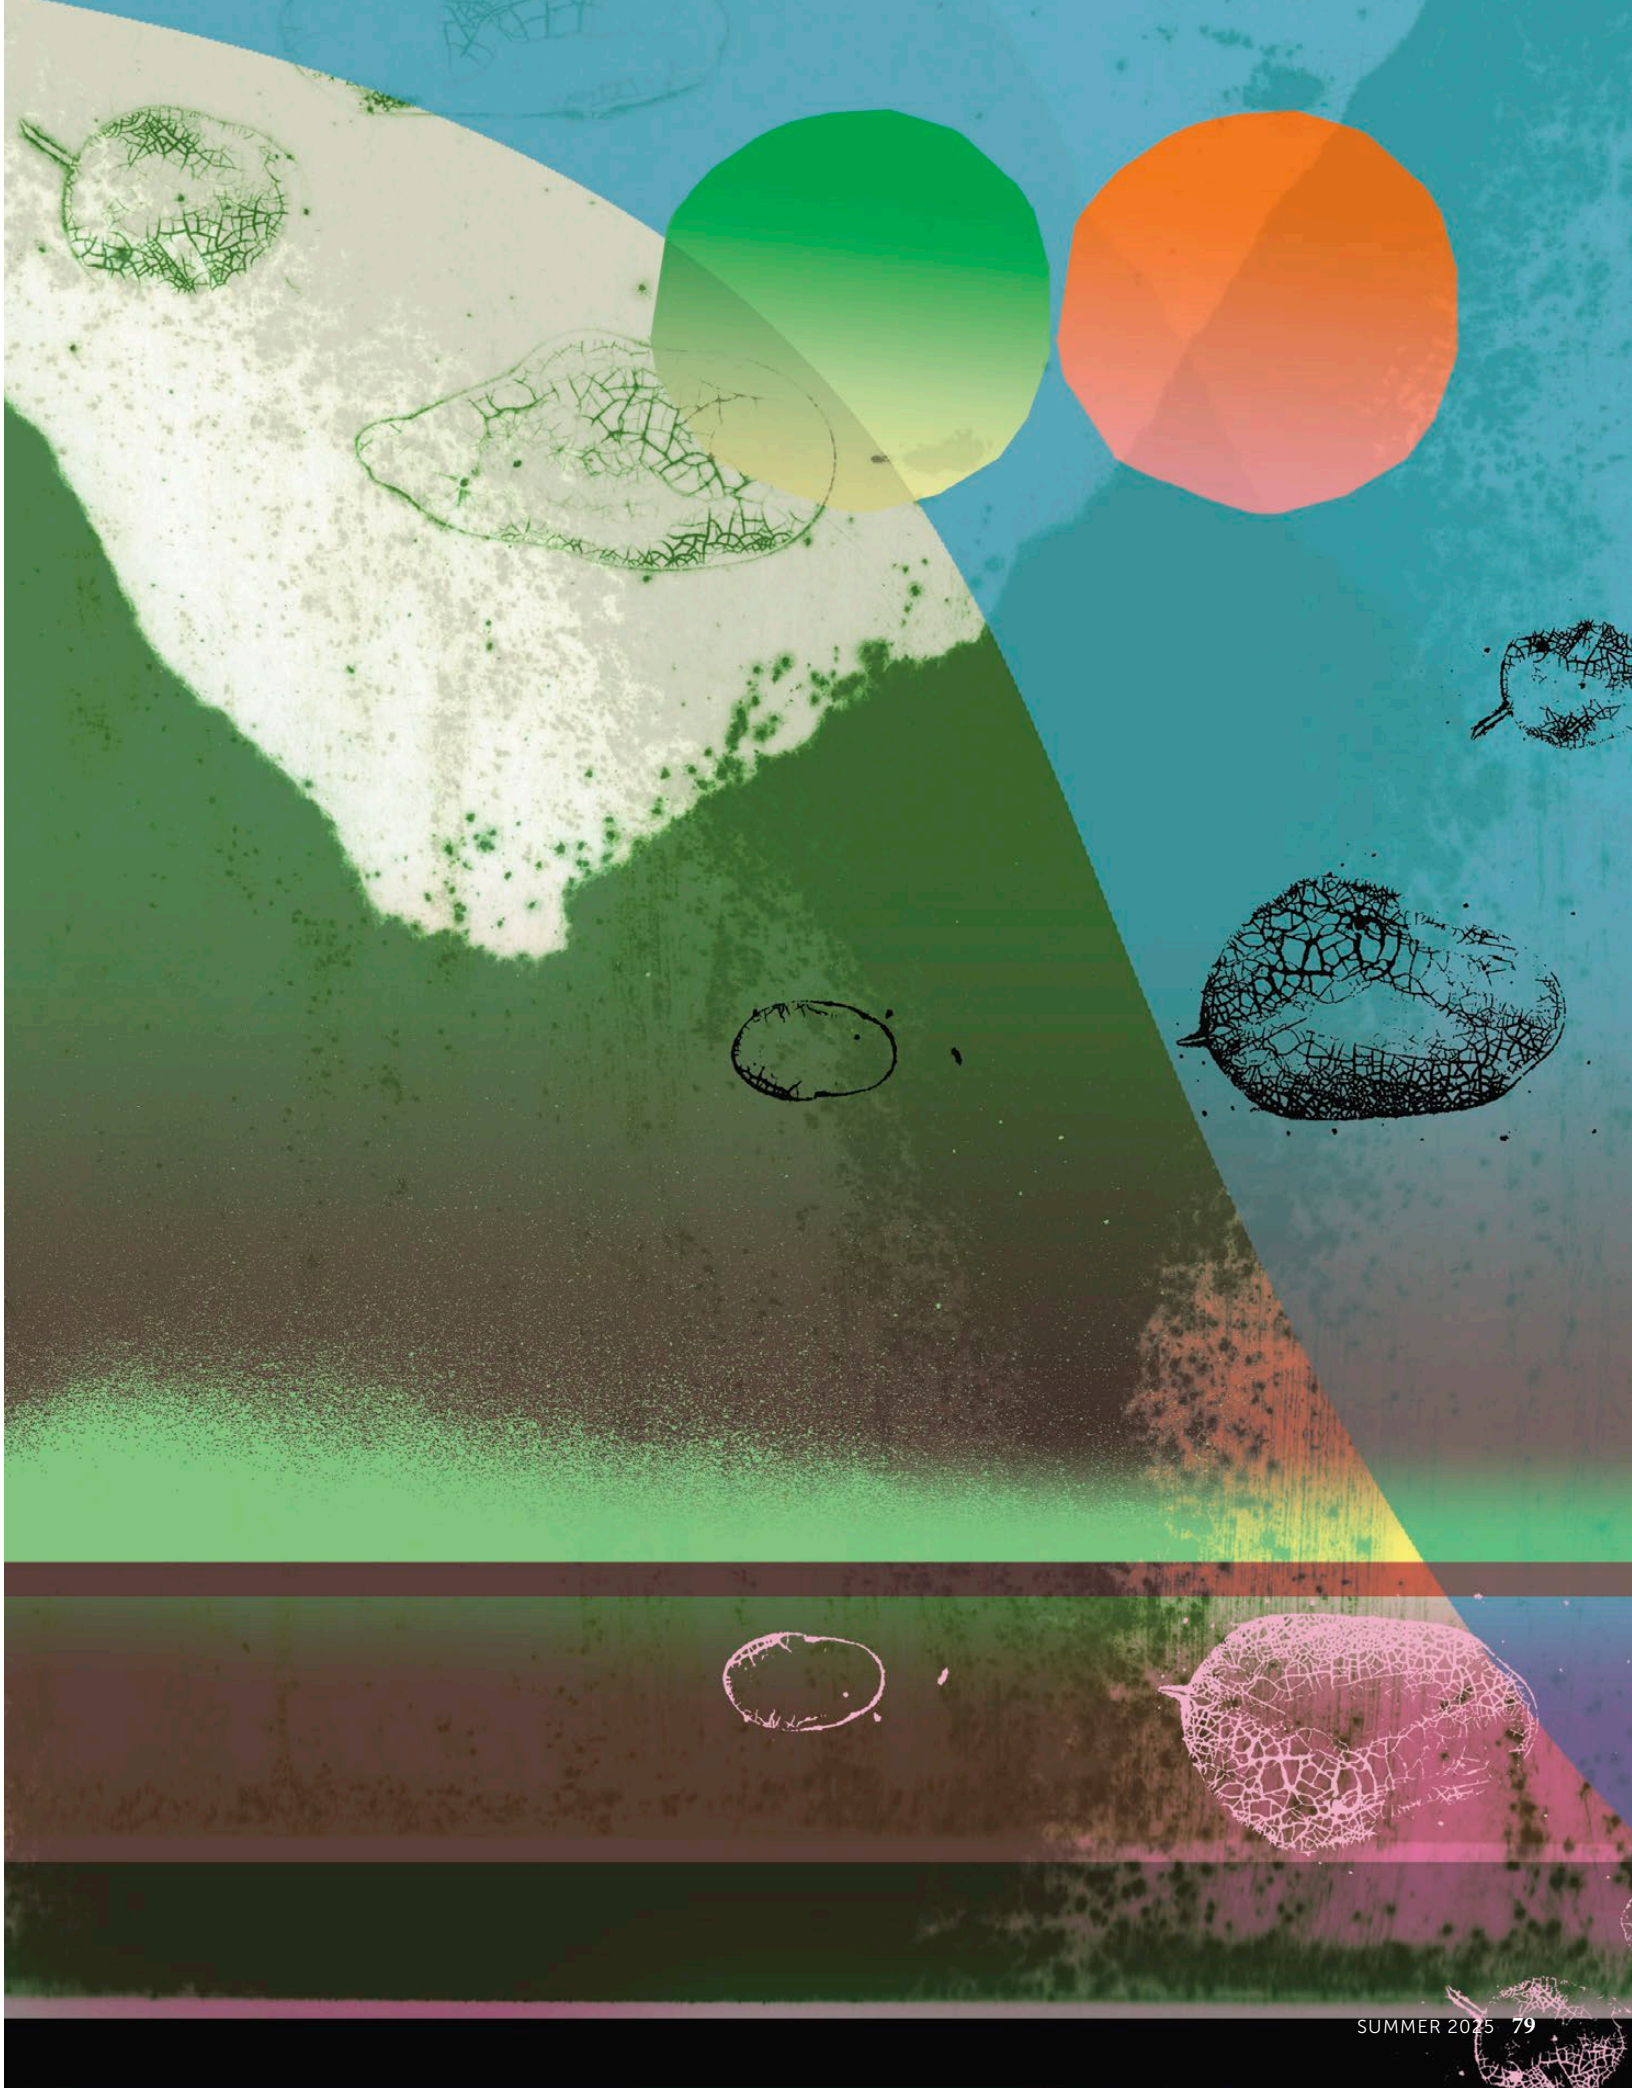

on a molecular level as well. The naturally occurring amino acids that make up proteins are designated as "left-handed," for example, and the nucleic acids in DNA and RNA are "right-handed." There are no known physical reasons why chiral biomolecules like amino acids and nucleic acids cannot exist in their opposite handed forms. In fact, scientists have assembled proteins from synthetic mirror-image amino acids, but those are not the forms that occur in nature. In a mirror life-form, cells would be identical to natural cells except the key biological molecules—nucleic acids, proteins, sugars, and lipids—would exist in chiral states opposite those found in nature.

Creating ordinary bacterial cells from synthetic molecular building blocks is currently impossible; doing so for mirror life would be more challenging, if it is even possible at all. However, many scientists have synthesized mirror proteins and other biomolecules. One lab has made mirror-image enzymes that can copy mirror-image DNA and has developed parts of a mirror ribosome, the molecular machine made of RNA that directs the synthesis of proteins.

It's worth noting that research on mirror biology could produce insights both practical and profound. Mirror biomolecules could spur a range of innovations, including new drugs that neither degrade nor cause immune reactions, as well as more efficient ways to manufacture medicines and materials. Studying mirror biomolecules could help scientists understand the origin of life and bring advances in cell replacement therapies, biomanufacturing, biosensing, and environmental remediation.

Studying mirror biomolecules could help scientists understand the origin of life and bring advances in cell replacement therapies, biomanufacturing, biosensing, and environmental remediation.

The statement in *Science* calls for more research to assess the risks and warns that without "compelling evidence for reassurance," no mirror bacteria should be created. It goes on to say, "We believe that this can be ensured with minimal impact on beneficial research." Yet the *Science* article says little about what lines should be drawn between acceptable research, unacceptable precursor research, and the prohibited creation of mirror bacteria. Without clear guidelines, the ban might be interpreted as allowing any precursor research, up to the point where making a mirror bacteria becomes almost trivial. Alternatively, it could be interpreted as not allowing any research that could someday be used to build mirror bacteria. That could stifle biochemistry and synthetic biology more broadly. To take this to an absurdist endpoint: All biochemistry produces precursor knowledge.

What's needed is a clear, practicable line that can be applied to relevant research and assure scientists, funders, and the public that work does not violate an ethical norm. An example of such a line can be found in human gene editing: The *somatic/germline divide* sets the stage such that research is allowed to modify cells in existing humans, so long as modifications cannot be passed to subsequent generations. We will not call for any particular line here. Instead, we will outline possibilities and a framework for considering them.

Reassessment points

The *precautionary principle* is often invoked when the potential outcome of research could lead to severe harm. There are several iterations of it, but the basic form is thus: If there are possible harms, and science's understanding of how those harms would come about is uncertain, then we shouldn't just barrel ahead. This makes intuitive sense: It's better to be safe than sorry. However, conceptually and practically, precaution is hard to get right. It frequently turns into—and is easily deployed as—a tool to bring science to a stop.

But when a technology has potential for benefit as well as harm, a permanent ban on its development could itself be harmful. So we argue that a successful precautionary approach should serve as a careful mediator between permission and prohibition. It might be summed up as "look before you leap." In 2016, a committee of the National Academies of Sciences, Engineering, and Medicine (NASEM) proposed an approach along these lines for gene drive research, essentially arguing that if there's a great benefit to leaping—or great harm in not leaping—the research should probably move forward, though as carefully as possible. One approving analysis of the report called it an "iterative approach to uncertainty."

Gene drives, a form of genetic engineering that ensures desired genetic modifications are passed on to descendants, demonstrate how this approach can work in practice. One potential use of gene drives is to combat malaria by eliminating disease-carrying mosquitos. The NASEM report proposed breaking research into phases to allow critical, constructive deliberation—starting with contained trials of genetically modified mosquitos before moving to field trials and then to post-release monitoring—with reassessments (and potential stopping points) built into each step. For human gene editing research, striking this balance between permission and

prohibition meant drawing a line between editing germline cells, which would create heritable changes, and editing somatic cells, which wouldn't, as well as creating a governance mechanism—via a designated National Institutes of Health committee—to assess proposals to modify somatic cells.

So how could "look before you leap" work in terms of mirror life? Creating mirror cells—if it can be done—will be a multistep process. Each step will have distinct benefits as well as possible risks. The field needs ways to pause and reassess research along that path—well before a mirror cell is created. The key is finding spots to stop the advance toward mirror life and create conceptual space to reassess. In other words, we need to look for places to draw lines.

Here, the concept of the slippery slope is helpful. In this metaphor, research that is unquestionably morally acceptable is at the top; the dystopian scenario that must be avoided sits at the bottom. A slope is slippery when reaching one point essentially foretells reaching the next point, meaning that once research steps off the top of the slope, where acceptable research lives, it will likely slip to undesirable points lower down. In the case of mirror life, basic biochemistry is at the top, and the release of a mirror bacteria into the world is at the bottom.

But this slide to the bottom doesn't have to be inevitable. What is needed are barriers along the slope; research above a given barrier would be acceptable and permitted, and below not. Moreover, each barrier would provide a place where scientists and society could stop and reassess whether the benefits and risks justify that barrier and, if not, to consider what the next barrier should be. In this scheme, the precautionary principle does not automatically become a ban on beneficial research but instead enables appropriately cautious research behavior.

An ineffective barrier is one that cannot clearly distinguish upslope research from downslope. One example comes from a NASEM report on human gene editing. It argued that editing for serious disease was acceptable (and upslope), whereas editing for nonserious disease was not (and downslope). But what does "serious" mean? Everyone can agree that sickle cell disease is serious, but what about deafness or extreme lack of height? Gene editing for sickle cell is already underway, so according to this framework, research can just slip down the slope of editing genetic traits of increasing ambiguity regarding their being "serious" or even "diseases."

Another kind of weak barrier is one that lacks moral weight. Imagine trying to draw a line on a slope concerning which animals can be made into human-animal chimeric embryos by designating a barrier based on animal intelligence. Assuming researchers could measure it, such intelligence will be a continuous scale, and there would be no justifiable reason for stopping research at rats instead of pigs, or whatever point on the slope is selected.

By contrast, the somatic/germline divide is an example of an effective barrier that has held, famously, for 50 years. Part of its durability lies in the fact that it is grounded in something anyone can grasp. Any modification upslope of the germline editing barrier applies to an existing individual, downslope to descendants or even the human species. Indeed, this distinction is now built into law in most European countries. Upslope has been portrayed as medicine, and downslope as not respecting "God's will" or "nature." The barrier is so clear that it cannot be crossed casually. When more precise gene editing technologies were invented that could be used to modify human embryos, and thus the germline, the barrier created the expectation that people stop and reassess.

We believe clear lines and strong barriers can be established for many, if not most, technologies to help think through and enact ethical controls. Barriers would ideally mark scientific, ontological, or cultural distinctions. We worry the *Science* article will convince readers that the slope around mirror life is so slippery, and the danger of hitting the bottom so high, that research should stop where it is. But the right sort of barriers can allow research to proceed and yield benefits, much as the barriers on the human gene editing slope have enabled breakthroughs like somatic treatments for sickle cell anemia.

One benefit of choosing to use barriers to navigate uncertainty is that they can be enforced in multiple ways, through law, policy, or social sanction. When getting laws passed is impractical or undesired, institutional policy and social norms can do much of the work of preventing harmful research because most researchers do not want to face formal discipline or be labeled a rogue scientist.

Mirror barriers

To spur discussion about mirror life, we propose five barrier candidates that could apply. These depend on biological distinctions in cells; other barriers might be based on moral or social distinctions. Early barriers concern the creation of cell components and later barriers, cell reproduction. In this way, the barriers are staged from more benign and feasible applications to more dangerous ones in which the requisite know-how remains undeveloped.

Creation of a non-mirror ribosome. Ribosomes, the macromolecular machine responsible for protein synthesis, are essential to a functioning cell, but cannot yet be created de novo in a lab. Efforts to make a mirror ribosome would certainly start by working out how to make a non-mirror ribosome from scratch. Setting the creation of a non-mirror ribosome as a barrier would arrest any further slide down the slope to making mirror life. However, stopping research at this barrier would also forfeit beneficial research in non-mirror biology, such as a more complete understanding of protein synthesis or production of more powerful protein drugs. As scientists approach this barrier, it will be important to stop and assess risks and benefits.

Creation of a mirror ribosome. The next downslope barrier, one less likely to impede non-mirror research, would be the construction of a mirror ribosome. Without a mirror ribosome to synthesize mirror proteins, there cannot be a viable cell, so creation of mirror versions of any other cell components could proceed without fear of slippage.

One reason the mirror ribosome is a strong barrier is that it would be crystal clear if a scientist had crossed it. That would be enough to pause research for further risk analysis. Such analysis would weigh potential benefits of having a mirror ribosome itself, such as synthesizing mirror pharmaceuticals that are more stable and less immunogenic. It would also allow consideration of other barriers that could suppress risks while enabling benefits. Given that the next barrier we can see on the slope is considerably closer to the creation of mirror life, creating a mirror ribosome could serve as the prohibited step for now.

Creation of a cell wall or membrane enclosing mirror machinery. Moving further down the slope, the next clear barrier is constructing a mirror cell wall or membrane that encapsulates mirror biomachinery. Since it is obvious when a cell wall is being created, there is a clear upslope and downslope. There is no cell without an enclosure, and the statement in *Science* does not claim that cell components are dangerous by themselves. This barrier also has ontological significance because there can be no distinct entity called a "cell" or "bacteria" without such enclosure. However, this barrier is quite close to the bottom of the slope: Were it crossed, the creation of mirror cells could be difficult to stop—especially if scientists or commercial interests were already invested in moving past this barrier.

Blocked reproduction in the lab. Notably, the *Science* article sees no danger in a single, isolated cell; the threat is in the potential for out-of-control growth. So, one possible barrier would be to ensure that any cell created would not be able to replicate. This has the advantage of being a clear boundary scientifically, as well as socially and culturally. Self-replicating entities are widely recognized as dangerous and are the stuff of dystopian television and movies. However, it is hard to imagine why scientists would engage in the required research to reach this barrier if they were prohibited from pressing on to make a cell that could replicate.

Blocked reproduction beyond the lab. At the bottom of the slope, according to the *Science* article, is the danger of mirror cells replicating outside of labs. Another way to create a barrier would be to ensure that any mirror bacteria synthesized be rendered incapable of living outside of very narrow, tightly controlled conditions.

Further down the slope, even below what the *Science* article posits as the bottom, would be creating a fully capable mirror bacteria but not releasing it. Though clear, this would be highly problematic as a barrier: It is so close to the unacceptable bottom of the slope that there is little to stop a rogue (or simply careless) actor from crossing it. Indeed, we think the entire point of the statement in *Science* is to safeguard us from getting anywhere close to this point.

Setting and shifting barriers

Barriers can both prevent a slide down the slippery slope and enable deliberate decisions on how to advance. If scientists and the public understand when work would cross a barrier, each barrier becomes a place to pause as knowledge accumulates and society assesses potential benefits and risks. What's more, the barrier framework enables questions to form higher up on the slope. For example, how valuable is the knowledge below the mirror ribosome barrier to science and the broader public?

Assessing this requires a social-scientific approach. Decisions about scientific technologies that might impact humanity should be based not only on the perspectives of scientific societies, foundations, or self-appointed groups of scientists and bioethicists, but also on the broader perspective of society—perhaps represented in a group that can make general and formal recommendations. Scientists and bioethicists who propose limits often make assumptions about what the public values. In contrast, barriers could be points to invite public deliberation.

The public's values should drive science policy at an abstract level, but there is also a clear need for expertise to connect such values to concrete action. After all, a regular member of the public does not have time to gain in-depth knowledge of the relevant biology, philosophy, or social science. Our group has shown the merit of scholars from different disciplinary backgrounds discussing the implications of technology development, and we hope that future discussions of mirror life are also interdisciplinary.

The statement in *Science* is an important warning. Scientists should consider the risks of creating mirror life. However, to heed that warning without foreclosing useful discoveries, there need to be guidelines about what precursor research is acceptable. We are not offering a solution, but we hope we have provided a framework for how to arrive at one.

John H. Evans *is Tata Chancellor's Chair of Social Sciences and codirector of the Institute for Practical Ethics at the University of California, San Diego (UCSD).* ***Craig Callender*** *is Tata Chancellor's Chair of Humanities and codirector of the Institute for Practical Ethics at UCSD.* ***Neal K. Devaraj*** *is a professor and Murray Goodman Endowed Chair of Chemistry and Biochemistry at UCSD.* ***Farren J. Isaacs*** *is a professor of molecular, cellular, and developmental biology and of biomedical engineering at Yale University.* ***Gregory E. Kaebnick*** *is a scholar at the Hastings Center who considers values at stake in synthetic biology and emerging genetic technologies.*