



United States Department of State

Washington, D.C. 20520

May 16, 2025

Case No. FL-2022-00062

Mr. Gary Ruskin
U.S. Right to Know
4096 Piedmont Avenue, #963
Oakland, CA 94611

Dear Mr. Ruskin:

As we noted in our letter dated April 4, 2025, we are processing your request for material under the Freedom of Information Act ("FOIA"), 5 U.S.C. § 552. The Department of State ("Department") has identified one additional responsive record subject to the FOIA. Upon review, we have determined this one record may be released in part.

An enclosure explains the FOIA exemptions and other grounds for withholding material. Where we have made redactions, the applicable FOIA exemptions are marked on the record. Where applicable, the Department has considered the foreseeable harm standard when reviewing these records and applying FOIA exemptions. All non-exempt material that is reasonably segregable from the exempt material has been released and is enclosed.

We will keep you informed as your case progresses. If you have any questions, your attorney may contact Assistant United States Attorney Stephanie Johnson at stephanie.johnson5@usdoj.gov or (202) 252-7874. Please refer to the case number, FL-2022-00062, and the civil action number, 22-cv-01130, in all correspondence about this case.

Sincerely,

A handwritten signature in cursive script that reads "Avery Bullard".

Avery Bullard
Supervisory Government Information Specialist
Office of Information Programs and Services

Enclosures: As stated.

Freedom of Information Act (5 U.S.C. § 552) and Privacy Act (5 U.S.C. § 552a)

FOIA Exemptions

- (b)(1) Information specifically authorized by an executive order to be kept secret in the interest of national defense or foreign policy. Executive Order 13526 includes the following classification categories:
- 1.4(a) Military plans, systems, or operations
 - 1.4(b) Foreign government information
 - 1.4(c) Intelligence activities, sources or methods, or cryptology
 - 1.4(d) Foreign relations or foreign activities of the US, including confidential sources
 - 1.4(e) Scientific, technological, or economic matters relating to national security, including defense against transnational terrorism
 - 1.4(f) U.S. Government programs for safeguarding nuclear materials or facilities
 - 1.4(g) Vulnerabilities or capabilities of systems, installations, infrastructures, projects, plans, or protection services relating to US national security, including defense against transnational terrorism
 - 1.4(h) Weapons of mass destruction
- (b)(2) Related solely to the internal personnel rules and practices of an agency
- (b)(3) Specifically exempted from disclosure by statute (other than 5 USC 552), for example:
- | | |
|----------------|---|
| ARMSEXP | Arms Export Control Act, 50a USC 2411(c) |
| CIA PERS/ORG | Central Intelligence Agency Act of 1949, 50 USC 403(g) |
| EXPORT CONTROL | Export Administration Act of 1979, 50 USC App. Sec. 2411(c) |
| FS ACT | Foreign Service Act of 1980, 22 USC 4004 |
| INA | Immigration and Nationality Act, 8 USC 1202(f), Sec. 222(f) |
| IRAN | Iran Claims Settlement Act, Public Law 99-99, Sec. 505 |
- (b)(4) Trade secrets and confidential commercial or financial information
- (b)(5) Interagency or intra-agency communications forming part of the deliberative process, attorney-client privilege, or attorney work product
- (b)(6) Personal privacy information
- (b)(7) Law enforcement information whose disclosure would:
- (A) interfere with enforcement proceedings
 - (B) deprive a person of a fair trial
 - (C) constitute an unwarranted invasion of personal privacy
 - (D) disclose confidential sources
 - (E) disclose investigation techniques
 - (F) endanger life or physical safety of an individual
- (b)(8) Prepared by or for a government agency regulating or supervising financial institutions
- (b)(9) Geological and geophysical information and data, including maps, concerning wells

Other Grounds for Withholding

- NR Material not responsive to a FOIA request excised with the agreement of the requester

| | |
|-----------------|--|
| From: | "Asher, David" (b)(6)@state.gov> |
| To: | Kanapathy, Ivan (b)(6) |
| CC: | Stilwell, David R (b)(6)@state.gov>; Keshap, Atul (b)(6)@state.gov>; Feith, David (b)(6)@state.gov>; DiNanno, Thomas G (b)(6)@state.gov>; Matthew Pottinger (b)(6); Gray, Alexander (b)(6); Hooker, Allison M (b)(6) |
| Subject: | SBU COVID Timeline v06-COVID origins "double opened " |
| Date: | Fri, 1 Jan 2021 17:57:05 +0000 |

This massive Excel spreadsheet timeline(200+ pages) is being shared just in case you are interested in what our research team has found hiding in plain sight on COV19 origins and PRC complicity going back many years. As someone who has been involved in USG criminal investigations against foreign adversaries for decades, the level of conspiracy evidenced in this unclassified spreadsheet is extreme. Of course, our even more detailed high side version is mind melting (if you don't have copy, I can send it on Monday). I am proud of what our little team has accomplished that no one else in the USG seems to have done to try to put the pieces of the puzzle together.

We also will have a very detailed scientific paper, informally commissioned by State, for circulation in the next couple days. It will address multiple questions on the genetic sequencing and will employ actual biostatistics to assess the likelihood of the natural zoonotic versus genetic gain of function hypotheses. It will be worth reading.

For the record, (b)(5) Deliberative Process

(b)(5) Deliberative Process

(b)(5)
Deliberati

See attached paper on the genetic structure of SARS COV 2.

(b)(5) Deliberative Process

I look forward to inviting all of you to discuss "what happened" and what we can do to prevent a future much worse syn-bio disaster at the Hudson Institute in the next month or two after Jan 20. Thanks for your remarkable, professional, and inspiring leadership and service. I am glad I came out of the shadows to do some productive work with a great team at State. No one will blame anyone on this email for not doing their very best to counter and protect the US from the ravages of COVID 19 as well as investigate the origins.

Happy New Year!

David

Some thoughts FWIW on a demarche:

(b)(5) Deliberative Process

If it's a question of timing or synchronization of the demarche, those are valid concerns. (b)(5)

(b)(5) Deliberative Process

From: (b)(6)@state.gov>
Sent: Monday, December 28, 2020 7:01 PM
To: DiNanno, Thomas G (b)(6)@state.gov>; Feith, David (b)(6)@state.gov>; Gross, Laura J (b)(6)@state.gov>; (b)(6)@state.gov>; (b)(6)@state.gov>; Gibbs, Jeffrey J (b)(6)@state.gov>; Paulopol, Andreea I (b)(6)@state.gov>
Cc: (b)(6)@state.gov>; (b)(6)@state.gov>; Asher, David (b)(6)@state.gov>
Subject: ~~SBU~~ COVID Timeline v06

Attached FYI: I have updated the SBU version of the timeline to format it to print on legal size paper with page numbering (all tabs). As previously mentioned to some, it contains 606 unique excerpts from 1985 to 11/11/2020. Along with the complete timeline (first/red tab), I have extracted several thematic timelines which you can find in the tabs to the right. In order left to right they are labeled:

- "Censorship of health info"
- "Delayed admitting human xmsn"
- "Limited, false, delayed reporting"
- "Exporting the virus"
- "Catastrophic missteps"
- "Efforts to counter lab hypthesis"
- "WHO as PRC cheerleader"
- "GOF research"
- "Handling lethal pathogens"
- "Poor safety, lab leak history"
- "US offers of support"

VR (b)(6)

PS: As of tomorrow, I will only be reachable via this email and my cell (b)(6) until the 6th.

(b)(6)

*Bureau of Arms Control, Verification and Compliance
US Department of State*

(b)(6)

NSTS (b)(6) NSTS)

JWICS (b)(6)@state.ic.gov

SIPR: (b)(6)@state.sgov.gov

~~SENSITIVE BUT UNCLASSIFIED~~

~~SENSITIVE BUT UNCLASSIFIED~~


Sender: "Asher, David" (b)(6)@state.gov>

Recipient:

Kanapathy, Ivan (b)(6)
Stilwell, David R (b)(6)@state.gov>;
Keshap, Atul (b)(6)@state.gov>;
Feith, David (b)(6)@state.gov>;
DiNanno, Thomas G (b)(6)@state.gov>;
Matthew Pottinger (b)(6)
Gray, Alexander (b)(6)
Hooker, Allison M (b)(6)

This PDF is available at <http://nap.edu/24890>

SHARE



Biodefense in the Age of Synthetic Biology (2018)

DETAILS

188 pages | 8.5 x 11 | PAPERBACK

ISBN 978-0-309-46518-2 | DOI 10.17226/24890

CONTRIBUTORS

Committee on Strategies for Identifying and Addressing Potential Biodefense Vulnerabilities Posed by Synthetic Biology; Board on Chemical Sciences and Technology; Board on Life Sciences; Division on Earth and Life Studies; National Academies of Sciences, Engineering, and Medicine

SUGGESTED CITATION

National Academies of Sciences, Engineering, and Medicine 2018. *Biodefense in the Age of Synthetic Biology*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/24890>.

GET THIS BOOK

FIND RELATED TITLES

Visit the National Academies Press at NAP.edu and login or register to get:

- Access to free PDF downloads of thousands of scientific reports
- 10% off the price of print titles
- Email or social media notifications of new titles related to your interests
- Special offers and discounts



Distribution, posting, or copying of this PDF is strictly prohibited without written permission of the National Academies Press. (Request Permission) Unless otherwise indicated, all materials in this PDF are copyrighted by the National Academy of Sciences.

Copyright © National Academy of Sciences. All rights reserved.

BIODEFENSE IN THE AGE OF Synthetic Biology

Committee on Strategies for Identifying and Addressing Potential Biodefense Vulnerabilities
Posed by Synthetic Biology

Board on Chemical Sciences and Technology

Board on Life Sciences

Division on Earth and Life Studies

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

THE NATIONAL ACADEMIES PRESS

Washington, DC

www.nap.edu

THE NATIONAL ACADEMIES PRESS**500 Fifth Street, NW****Washington, DC 20001**

This project was supported by Contract No. HQ0034-16-C-0062 between the National Academy of Sciences and the U.S. Department of Defense. Any opinions, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect the view of any organization or agency that provided support for the project.

International Standard Book Number-13: 978-0-309-46518-2

International Standard Book Number-10: 0-309-46518-4

Library of Congress Control Number: 2018911261

Digital Object Identifier: <https://doi.org/10.17226/24890>

Copyright 2018 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

Suggested citation: National Academies of Sciences, Engineering, and Medicine. 2018. *Biodefense in the Age of Synthetic Biology*. Washington, DC: The National Academies Press. doi: <https://doi.org/10.17226/24890>.

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

The **National Academy of Sciences** was established in 1863 by an Act of Congress, signed by President Lincoln, as a private, nongovernmental institution to advise the nation on issues related to science and technology. Members are elected by their peers for outstanding contributions to research. Dr. Marcia McNutt is president.

The **National Academy of Engineering** was established in 1964 under the charter of the National Academy of Sciences to bring the practices of engineering to advising the nation. Members are elected by their peers for extraordinary contributions to engineering. Dr. C. D. Mote, Jr., is president.

The **National Academy of Medicine** (formerly the Institute of Medicine) was established in 1970 under the charter of the National Academy of Sciences to advise the nation on medical and health issues. Members are elected by their peers for distinguished contributions to medicine and health. Dr. Victor J. Dzau is president.

The three Academies work together as the **National Academies of Sciences, Engineering, and Medicine** to provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The National Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

Learn more about the National Academies of Sciences, Engineering, and Medicine at www.nationalacademies.org.

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

Consensus Study Reports published by the National Academies of Sciences, Engineering, and Medicine document the evidence-based consensus on the study's statement of task by an authoring committee of experts. Reports typically include findings, conclusions, and recommendations based on information gathered by the committee and the committee's deliberations. Each report has been subjected to a rigorous and independent peer-review process and it represents the position of the National Academies on the statement of task.

Proceedings published by the National Academies of Sciences, Engineering, and Medicine chronicle the presentations and discussions at a workshop, symposium, or other event convened by the National Academies. The statements and opinions contained in proceedings are those of the participants and are not endorsed by other participants, the planning committee, or the National Academies.

For information about other products and activities of the National Academies, please visit www.nationalacademies.org/about/whatwedo.

**COMMITTEE ON STRATEGIES FOR IDENTIFYING AND ADDRESSING POTENTIAL
BIODEFENSE VULNERABILITIES POSED BY SYNTHETIC BIOLOGY**

Members

MICHAEL IMPERIALE (*Chair*), University of Michigan Medical School
PATRICK BOYLE,¹ Ginkgo Bioworks
PETER A. CARR, Massachusetts Institute of Technology Lincoln Laboratory
DOUGLAS DENSMORE, Boston University
DIANE DIEULIIS, National Defense University
ANDREW ELLINGTON, University of Texas at Austin
GIGI KWIK GRONVALL, Johns Hopkins Center for Health Security
CHARLES HAAS, Drexel University
JOSEPH KANABROCKI, University of Chicago
KARA MORGAN, Quant Policy Strategies, LLC
KRISTALA JONES PRATHER, Massachusetts Institute of Technology
THOMAS SLEZAK, Lawrence Livermore National Laboratory
JILL TAYLOR, New York State Department of Health

Staff

MARILEE SHELTON-DAVENPORT, Study Director
KATHERINE BOWMAN, Senior Program Officer
JENNA OGILVIE, Research Associate
JARRETT NGUYEN, Senior Program Assistant

¹ See Appendix E, Disclosure of Conflict of Interest.

BOARD ON CHEMICAL SCIENCES AND TECHNOLOGY*Co-Chairs***DAVID BEM**, PPG Industries**JOAN BRENNECKE**, NAE, University of Texas, Austin*Members***GERARD BAILLELY**, Procter and Gamble**MARK A. BARTEAU**, NAE, Texas A&M**MICHELLE V. BUCHANAN**, Oak Ridge National Laboratory**JENNIFER SINCLAIR CURTIS**, University of California, Davis**RICHARD EISENBERG**, NAS, University of Rochester**SAMUEL H. GELLMAN**, NAS, University of Wisconsin–Madison**SHARON C. GLOTZER**, NAS, University of Michigan**MIRIAM E. JOHN**, Sandia National Laboratories (*retired*)**ALAN D. PALKOWITZ**, Eli Lilly and Company (*retired*)**JOSEPH B. POWELL**, Shell**PETER J. ROSSKY**, NAS, Rice University**RICHMOND SARPONG**, University of California, Berkeley**TIMOTHY SWAGER**, NAS, Massachusetts Institute of Technology*National Academies of Sciences, Engineering, and Medicine Staff***TERESA FRYBERGER**, Board Director**MARILEE SHELTON-DAVENPORT**, Senior Program Officer**CAMLY TRAN**, Senior Program Officer**ANNA SBEREGAEVA**, Associate Program Officer**JARRETT I. NGUYEN**, Senior Program Assistant**JESSICA WOLFMAN**, Senior Program Assistant**SHUBHA BANSKOTA**, Financial Associate

BOARD ON LIFE SCIENCES*Chair*

JAMES P. COLLINS, Arizona State University

Members

ALONSO AGUIRRE, George Mason University

ENRIQUETA C. BOND, Burroughs Wellcome Fund

DOMINIQUE BROSSARD, University of Wisconsin-Madison

ROGER D. CONE, University of Michigan

NANCY D. CONNELL, Rutgers New Jersey Medical School

SEAN M. DECATUR, Kenyon College

JOSEPH R. ECKER, Salk Institute for Biological Studies

SCOTT V. EDWARDS, Harvard University

GERALD L. EPSTEIN, National Defense University

ROBERT J. FULL, University of California, Berkeley

ELIZABETH HEITMAN, Vanderbilt University Medical Center

JUDITH KIMBLE, University of Wisconsin–Madison

MARY E. MAXON, Lawrence Berkeley National Laboratory

ROBERT NEWMAN, Independent Consultant

STEPHEN J. O'BRIEN, Nova Southeastern University

CLAIRE POMEROY, Albert and Mary Lasker Foundation

MARY E. POWER, University of California, Berkeley

SUSAN RUNDELL SINGER, University of California, Berkeley

LANA SKIRBOLL, Sanofi

DAVID R. WALT, Harvard Medical School

Staff

FRANCES SHARPLES, Director

KATHERINE BOWMAN, Senior Program Officer

ANDREA HODGSON, Program Officer

JO HUSBANDS, Senior Scholar

KEEGAN SAWYER, Senior Program Officer

AUDREY THEVENON, Program Officer

JESSICA DE MOUY, Senior Program Assistant

KOSSANA YOUNG, Senior Program Assistant

Acknowledgment of Reviewers

This Consensus Study Report was reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise. The purpose of this independent review is to provide candid and critical comments that will assist the National Academies of Sciences, Engineering, and Medicine in making each published report as sound as possible and to ensure that it meets the institutional standards for quality, objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process.

We thank the following individuals for their review of this report:

JAMES BURNS, Casebia Therapeutics
MICHAEL DIAMOND, Washington University School of Medicine
JAMES DIGGANS, Twist Bioscience
DREW ENDY, Stanford University
CAROLINE GENCO, Tufts University School of Medicine
JACQUELINE GIBSON, University of North Carolina, Chapel Hill
KAREN E. JENNI, U.S. Geological Survey
MICHAEL JEWETT, Northwestern University
GREGORY KAEBNICK, The Hastings Center
MARGARET E. KOSAL, Georgia Institute of Technology
KAREN E. NELSON, J. Craig Venter Institute
MICHAEL OSTERHOLM, University of Minnesota
TARA O'TOOLE, In-Q-Tel
PAMELA A. SILVER, Harvard Medical School
DAVID WALT, Harvard Medical School and Harvard University

Although the reviewers listed above provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations of this report nor did they see the final draft before its release. The review of this report was overseen by **MICHAEL LADISCH**, Purdue University, and **RANDALL MURCH**, Virginia Polytechnic and State University. They were responsible for making certain that an independent examination of this report was carried out in accordance with the standards of the National Academies and that all review comments were carefully considered. Responsibility for the final content rests entirely with the authoring committee and the National Academies.

Contents

| | |
|--|----|
| SUMMARY | 1 |
| 1 INTRODUCTION | 9 |
| Understanding Synthetic Biology, 11 | |
| Assessing Potential Biodefense Concerns, 12 | |
| Mitigating Potential Biodefense Concerns, 12 | |
| Study Approach, 13 | |
| 2 BIOTECHNOLOGY IN THE AGE OF SYNTHETIC BIOLOGY | 15 |
| What Is Synthetic Biology, 15 | |
| Implications of the Age of Synthetic Biology, 18 | |
| Specific Synthetic Biology Technologies and Applications, 21 | |
| 3 FRAMEWORK FOR ASSESSING CONCERN ABOUT SYNTHETIC BIOLOGY CAPABILITIES | 23 |
| Approach to Developing the Framework, 23 | |
| Factors for Assessing Concern, 25 | |
| Applying the Framework in the Assessment of Concern, 31 | |
| 4 ASSESSMENT OF CONCERNS RELATED TO PATHOGENS | 37 |
| Re-Creating Known Pathogens, 38 | |
| Making Existing Pathogens More Dangerous, 44 | |
| Creating New Pathogens, 54 | |
| Summary, 57 | |
| 5 ASSESSMENT OF CONCERNS RELATED TO PRODUCTION OF CHEMICALS OR BIOCHEMICALS | 59 |
| Manufacturing Chemicals or Biochemicals by Exploiting Natural Metabolic Pathways, 61 | |
| Manufacturing Chemicals or Biochemicals by Creating Novel Metabolic Pathways, 64 | |

| | |
|---|-----|
| Making Biochemicals via In Situ Synthesis, 65 | |
| Summary, 69 | |
| 6 ASSESSMENT OF CONCERNS RELATED TO BIOWEAPONS THAT ALTER THE HUMAN HOST | 71 |
| Modifying the Human Microbiome, 71 | |
| Modifying the Human Immune System, 74 | |
| Modifying the Human Genome, 77 | |
| Summary, 81 | |
| 7 RELATED DEVELOPMENTS THAT MAY IMPACT THE ABILITY TO EFFECT AN ATTACK USING A SYNTHETIC BIOLOGY-ENABLED WEAPON | 85 |
| Barriers to the Use of Bioweapons, 85 | |
| Relevant Convergent Technologies, 87 | |
| Summary, 92 | |
| 8 OPTIONS FOR MITIGATING CONCERNS | 95 |
| Current Mitigation Approaches and Infrastructure, 95 | |
| Mitigation Challenges Posed by Synthetic Biology, 100 | |
| Potential Opportunities to Advance Mitigation Capabilities, 104 | |
| Summary, 112 | |
| 9 MOVING FORWARD: CONCLUSIONS AND RECOMMENDATIONS | 115 |
| Overarching Recommendation, 116 | |
| Concerns Posed by Synthetic Biology-Enabled Capabilities, 116 | |
| Future Use of the Framework, 125 | |
| Biodefense Implications of the Age of Synthetic Biology, 126 | |
| REFERENCES | 129 |
| APPENDIXES | |
| A Specific Synthetic Biology Concepts, Approaches, and Tools | 147 |
| B Selected Prior Analyses Used to Inform the Framework | 157 |
| C Questions to Stimulate Consideration of Framework Factors | 163 |
| D Committee Biographies | 167 |
| E Disclosure of Conflict of Interest | 171 |
| F Study Methods | 173 |

Summary

Scientific advances over the past several decades have accelerated the ability to engineer existing organisms and to potentially create novel ones not found in nature. Synthetic biology, which collectively refers to concepts, approaches, and tools that enable the modification or creation of biological organisms, is being pursued overwhelmingly for beneficial purposes ranging from reducing the burden of disease to improving agricultural yields to remediating pollution. Although the contributions synthetic biology can make in these and other areas hold great promise, it is also possible to imagine malicious uses that could threaten U.S. citizens and military personnel. Making informed decisions about how to address such concerns requires a realistic assessment of the capabilities that could be misused. To that end, the U.S. Department of Defense, working with other agencies involved in biodefense, asked the National Academies of Sciences, Engineering, and Medicine to develop a framework to guide an assessment of the security concerns related to advances in synthetic biology, to assess the levels of concern warranted for such advances, and to identify options that could help mitigate those concerns. An excerpted version of the study charge highlights the key tasks undertaken (see Chapter 1, Box 1-2 for the more detailed statement of task):

To assist the U.S. Department of Defense's Chemical and Biological Defense Program (CBDP), the National Academies of Sciences, Engineering, and Medicine will appoint an ad hoc committee to address the changing nature of the biodefense threat in the age of synthetic biology. Specifically, the focus of the study will be the manipulation of biological functions, systems, or microorganisms resulting in the production of disease-causing agents or toxins. . . . Initially, the committee will develop a strategic framework to guide an assessment of the potential security vulnerabilities related to advances in biology and biotechnology, with a particular emphasis on synthetic biology.

The framework will focus on how to address the following three questions: What are the possible security concerns with regard to synthetic biology that are on the horizon? What are the time frames of development of these concerns? What are our options for mitigating these potential concerns? . . .

. . . [T]he committee will use the outlined strategic framework to generate an assessment of potential vulnerabilities posed by synthetic biology. Inputs to this assessment may include information about the current threat, current program priorities and research, and an evaluation of the current landscape of science and technology. Conclusions and recommendations will include a list and description of potential vulnerabilities posed by synthetic biology.

An initial framework for assessing concerns was published in an interim report (National Academies of Sciences, Engineering, and Medicine, 2017a). This, the study's final report, builds on and supersedes that report. This report

explores and envisions potential misuses of synthetic biology, including concepts that are regularly discussed in open meetings. The potential misuses as they are discussed in the report are neither comprehensive nor enabling in the level of information and detail provided; they are included to illustrate the expanding mission of biodefense in the age of synthetic biology.

OVERARCHING RECOMMENDATION

Biotechnology in the age of synthetic biology expands the landscape of potential defense concerns. The U.S. Department of Defense (DoD) and its partnering agencies should continue to pursue ongoing strategies for chemical and biological defense; these strategies remain relevant in the age of synthetic biology. DoD and its partners also need to have approaches to account for the broader capabilities enabled by synthetic biology, now and into the future.

The nation's experience preparing for naturally occurring diseases provides a strong foundation for developing strategies to prevent and respond to emerging biologically enabled threats, particularly those based on naturally occurring pathogens. But synthetic biology approaches also have the potential to be used in ways that could change the presentation of an attack, for example, by modifying the properties of existing microorganisms, using microorganisms to produce chemicals, or employing novel or unexpected strategies to cause harm. It is valuable for the U.S. government to pay close attention to rapidly advancing fields such as synthetic biology, just as it did to advances in chemistry and physics during the Cold War era. However, approaches modeled after those taken to counter Cold War threats are not sufficient to address biological and biologically enabled chemical weapons in the age of synthetic biology. The partners involved in the U.S. biodefense enterprise will need expanded strategies and approaches to account for the new capabilities enabled by advances in this field.

A FRAMEWORK FOR ASSESSING CONCERN CONTRIBUTES TO PLANNING

Recommendation

The Department of Defense and its interagency partners should use a framework in assessing synthetic biology capabilities and their implications.

- (a) **A framework is a valuable tool for parsing the changing biotechnology landscape.**
- (b) **Using a framework facilitates the identification of bottlenecks and barriers, as well as efforts to monitor advances in technology and knowledge that change what is possible.**
- (c) **A framework provides a mechanism for incorporating the necessary technical expertise into the assessment.** A framework enables the participation of technical experts in synthetic biology and biotechnology along with experts in complementary areas (e.g., intelligence and public health).

The framework developed in the report identifies the features of a synthetic biology-enabled capability that would increase or decrease the level of concern about a given capability being used for harm. As summarized in Figure S-1, this framework identifies factors to determine the relative levels of concern posed by advances in biotechnology. In addition to supporting the analysis conducted in this study, the framework is intended to aid others in their consideration of current and future synthetic biology capabilities. Specifically, the framework is designed to support uses including analyzing existing biotechnologies to evaluate the levels of concern warranted at present; understanding how various technologies or capabilities compare to, interact with, or complement each other; identifying key bottlenecks and barriers that, if removed, could lead to a change in the level of concern about a capability; evaluating the implications of new experimental results or new technologies; and horizon-scanning to predict or prepare for potential future areas of concern. Use of a framework for assessing the implications of

SUMMARY

3

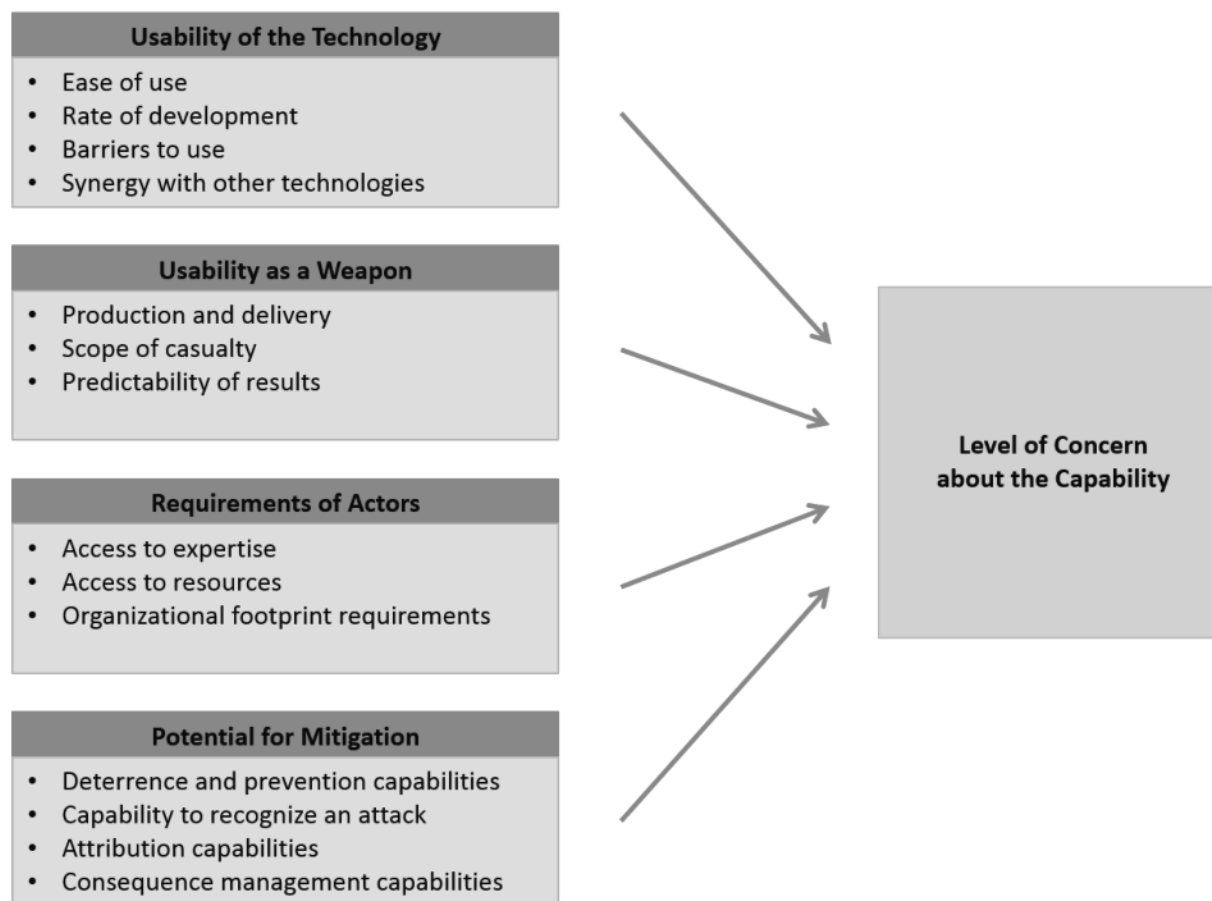


FIGURE S-1 Framework for assessing concern. The framework consists of four factors, along with descriptive elements within each factor. The factors are Usability of the Technology, Usability as a Weapon, Requirements of Actors, and Potential for Mitigation. These factors delineate the information used to assess the level of concern for particular synthetic biology-enabled capabilities.

synthetic biology capabilities thus contributes to biodefense planning and facilitates consideration of expert opinions about specific synthetic biology-enabled capabilities or combinations of capabilities.

SYNTHETIC BIOLOGY EXPANDS WHAT IS POSSIBLE

Synthetic biology expands what is possible in creating new weapons. It also expands the range of actors who could undertake such efforts and decreases the time required. Based on this study's analysis of the potential ways in which synthetic biology approaches and tools may be misused to cause harm, the following specific observations were made:

- Of the potential capabilities assessed, three currently warrant the most concern: (1) re-creating known pathogenic viruses, (2) making existing bacteria more dangerous, and (3) making harmful biochemicals via in situ synthesis.** The first two capabilities are of high concern due to usability of the

- technology. The third capability, which involves using microbes to produce harmful biochemicals in humans, is of high concern because its novelty challenges potential mitigation options.
- (b) **With regard to *pathogens*, synthetic biology is expected to (1) expand the range of what could be produced, including making bacteria and viruses more harmful; (2) decrease the amount of time required to engineer such organisms; and (3) expand the range of actors who could undertake such efforts.** The creation and manipulation of pathogens is facilitated by increasingly accessible technologies and starting materials, including DNA sequences in public databases. A wide range of pathogen characteristics could be explored as part of such efforts.
 - (c) **With regard to *chemicals, biochemicals, and toxins*, synthetic biology blurs the line between chemical and biological weapons.** High-potency molecules that can be produced through simple genetic pathways are of greatest concern, because they could conceivably be developed with modest resources and organizational footprint.
 - (d) **It may be possible to use synthetic biology to *modulate human physiology in novel ways*.** These ways include physiological changes that differ from the typical effects of known pathogens and chemical agents. Synthetic biology expands the landscape by potentially allowing the delivery of biochemicals by a biological agent and by potentially allowing the engineering of the microbiome or immune system. Although unlikely today, these types of manipulations may become more feasible as knowledge of complex systems, such as the immune system and microbiome, grows.
 - (e) **Some malicious applications of synthetic biology may not seem plausible now but could become achievable if certain barriers are overcome.** These barriers include knowledge barriers, as is the case for building a novel pathogen, or technological barriers, as in engineering complex biosynthetic pathways into bacteria or re-creating known bacterial pathogens. It is important to continue to monitor advances in biotechnology that may lower these barriers.

Synthetic biology concepts, approaches, and tools do not, in and of themselves, pose inherent harm. Rather, concerns derive from the specific applications or capabilities that synthetic biology might enable. The framework developed in the report was applied to assess the relative levels of concern posed by a set of synthetic biology capabilities. This assessment was undertaken in several steps. First, the framework was used to qualitatively analyze each of the identified capabilities individually. This analysis included considerations related to the state of the art of the technologies involved, the feasibility of using the capability to produce an effective weapon, the characteristics and resources an actor would likely require to carry out an attack, and information on proactive and reactive measures that might be taken to help mitigate the effects of misusing the capability. Then, an overall level of concern was determined for each capability relative to the other capabilities considered and an assessment of the landscape of capabilities and concerns presented. The results of this assessment are summarized in Figure S-2.

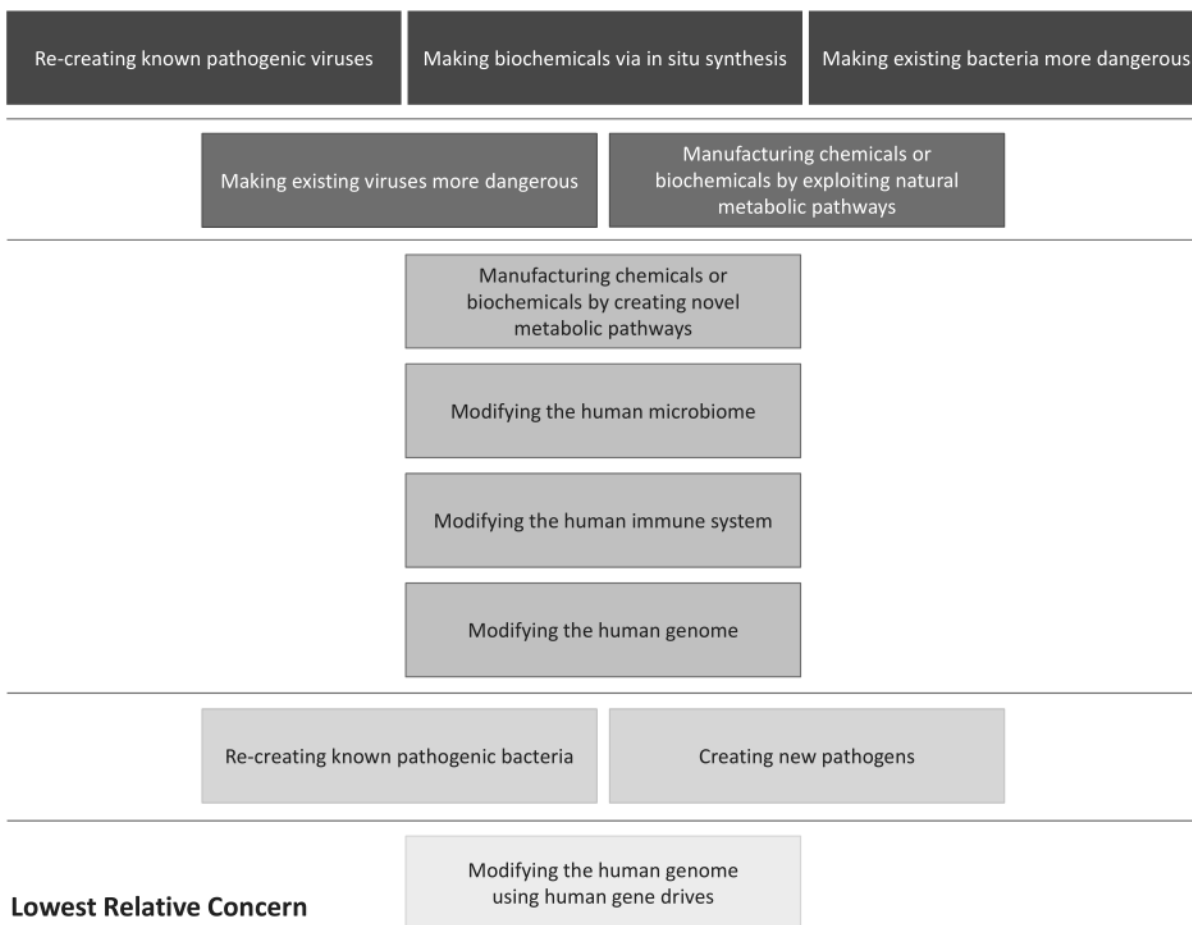
Capabilities currently warranting the highest relative level of concern include re-creating known pathogenic viruses, making biochemical compounds via *in situ* synthesis, and the use of synthetic biology to make existing bacteria more dangerous. These capabilities are based on technologies and knowledge that are readily available to a wide array of actors. Capabilities posing a moderate-to-high relative level of concern include manufacturing chemicals or biochemicals by exploiting natural metabolic pathways and the use of synthetic biology to make existing viruses more dangerous. These capabilities are also supported by available technologies and knowledge but involve more constraints and would likely be limited by factors related to both biology and skill. Capabilities posing a moderate relative level of concern include manufacturing chemicals or biochemicals by creating novel metabolic pathways, efforts to modify the human microbiome to cause harm, efforts to modify the human immune system, and efforts to modify the human genome. Although conceivable, these capabilities are more futuristic and likely limited by available knowledge and technology. Capabilities warranting a lower relative level of concern include re-creating known pathogenic bacteria and creating new pathogens; these capabilities involve major design and implementation challenges. The use of human gene drives warrants a minimal level of concern because it would be impractical to rely on generations of sexual reproduction to spread a harmful trait through a human population.

The application of the report's framework in this analysis reflects a snapshot in time, given understanding of current technologies and capabilities. As the field continues to evolve, some bottlenecks will likely widen and

SUMMARY

5

Highest Relative Concern



Lowest Relative Concern

FIGURE S-2 Relative ranking of concerns related to the synthetic biology-enabled capabilities analyzed. At present, capabilities toward the top warrant a relatively higher level of concern while capabilities toward the bottom warrant a relatively lower level of concern.

some barriers will be overcome. Table S-1 identifies a number of technical developments that may contribute to overcoming such bottlenecks and barriers to increase the feasibility or impact of a potential attack and the level of biodefense concern warranted for a capability. It is impossible to predict precisely when these developments might occur; those time lines are influenced by the drivers of commercial development and academic research, as well as by converging or synergistic technologies that may come from outside the field of synthetic biology. It will be important to continue to monitor advances in synthetic biology and biotechnology that may affect these bottlenecks and barriers.

TABLE S-1 Bottlenecks and Barriers That Currently Constrain the Capabilities Considered and Developments That Could Reduce These Constraints

| Capability | Bottleneck or Barrier | Relevant Developments to Monitor |
|--|---|--|
| Re-creating known pathogenic viruses | Bootling | Demonstrations of bootling viruses with synthesized genomes |
| Re-creating known pathogenic bacteria | DNA synthesis and assembly | Improvements in synthesis and assembly technology for handling larger DNA constructs |
| | Bootling | Demonstrations of bootling bacteria with synthesized genomes |
| Making existing viruses more dangerous | Constraints on viral genome organization | Increased knowledge of viral genome organization and/or demonstration of combinatorial approaches capable of facilitating larger-scale modifications to viral genome |
| | Engineering complex viral traits | Increased knowledge of determinants of complex viral traits, as well as how to engineer pathways to produce them |
| Making existing bacteria more dangerous | Engineering complex bacterial traits | Advances in combinatorial approaches and/or increased knowledge of determinants of complex bacterial traits, as well as how to engineer pathways to produce them |
| Creating new pathogens | Limited knowledge regarding minimal requirements for viability (in both viruses and bacteria) | Increased knowledge of requirements for viability in viruses or bacteria |
| | Constraints on viral genome organization | Increased knowledge of viral genome organization and/or demonstration of combinatorial approaches capable of facilitating larger-scale modifications to viral genome |
| Manufacturing chemicals or biochemicals by exploiting natural metabolic pathways | Tolerability of toxins to the host organism synthesizing the toxin | Pathway elucidation, improvements in circuit design, and improvements in host ("chassis") engineering to make toxins tolerable to the host organism synthesizing the toxin |
| | Pathway not known | Pathway elucidation and/or demonstrations of combinatorial approaches |
| | Challenges to large-scale production | Improvements in intracellular and industrial productivity |
| Manufacturing chemicals or biochemicals by creating novel metabolic pathways | Tolerability of toxins to the host organism synthesizing the toxin | Pathway elucidation and/or improvements in circuit design and/or improvements in host ("chassis") engineering to make toxins tolerable to the host organism synthesizing the toxin |
| | Engineering enzyme activity | Increased knowledge of how to modify enzymatic functions to make specific products |
| | Limited knowledge of requirements for designing novel pathways | Improvements in directed evolution and/or increased knowledge of how to build pathways from disparate organisms |
| | Challenges to large-scale production | Improvements in intracellular and industrial productivity |
| Making biochemicals via in situ synthesis | Limited understanding of microbiome | Improvements in knowledge related to microbiome colonization of host, in situ horizontal transfer of genetic elements, and other relationships between microbiome organisms and host processes |
| Modifying the human microbiome | Limited understanding of microbiome | Improvements in knowledge related to microbiome colonization of host, in situ horizontal transfer of genetic elements, and other relationships between microbiome organisms and host processes |

TABLE S-1 Continued

| Capability | Bottleneck or Barrier | Relevant Developments to Monitor |
|-----------------------------------|---|--|
| Modifying the human immune system | Engineering of delivery system | Increased knowledge related to the potential for viruses or microbes to deliver immunomodulatory factors |
| | Limited understanding of complex immune processes | Knowledge related to how to manipulate the immune system, including how to cause autoimmunity and predictability across a population |
| Modifying the human genome | Means to engineer horizontal transfer | Increased knowledge of techniques to effectively alter the human genome through horizontal transfer of genetic information |
| | Lack of knowledge about regulation of human gene expression | Increased knowledge related to regulation of human gene expression |

NOTE: Shading indicates developments thought to be propelled by commercial drivers. Some approaches, such as combinatorial approaches and directed evolution, may allow bottlenecks and barriers to be widened or overcome with less explicit knowledge or tools.

A RANGE OF STRATEGIES IS NEEDED TO PREPARE AND RESPOND

Recommendations

Many of the traditional approaches to biological and chemical defense preparedness will be relevant to synthetic biology, but synthetic biology will also present new challenges. The Department of Defense (DoD) and partner agencies will need approaches to biological and chemical weapons defense that meet these new challenges.

- (a) **The DoD and its partners in the chemical and biological defense enterprise should continue exploring strategies that are applicable to a wide range of chemical and biodefense threats.** Nimble biological and chemical defense strategies are needed because of rapid rates of technological change, as well as strategies adaptable to a wide range of threats because of uncertainty about which approaches an adversary might pursue.
- (b) **The potential unpredictability related to how a synthetic biology-enabled weapon could manifest creates an added challenge to monitoring and detection. The DoD and its partners should evaluate the national military and civilian infrastructure that informs population-based surveillance, identification, and notification of both natural and purposeful health threats.** An evaluation should consider whether and how the public health infrastructure needs to be strengthened to adequately recognize a synthetic biology-enabled attack. Ongoing evaluation will support responsive and adaptive management as technology advances.
- (c) **The U.S. government, in conjunction with the scientific community, should consider strategies that manage emerging risk better than current agent-based lists and access control approaches.** Strategies based on lists, such as the Federal Select Agent Program Select Agents and Toxins list, will be insufficient for managing risks arising from the application of synthetic biology. While measures to control access to physical materials such as synthetic nucleic acids and microbial strains have merits, such approaches will not be effective in mitigating all types of synthetic biology-enabled attacks.

Exploration Areas

It has been stated by both scientific and political leaders that the 21st century is the century of the life sciences. But as with previous expansions in technological capabilities, biotechnology in the age of synthetic biology presents a “dual-use dilemma” that scientific knowledge, materials, and techniques required for beneficial research or development could be misused to cause harm. Although current approaches to defense and public health preparedness remain valuable, there are also clear limitations to current approaches such as pathogen list–based screening tools.

To comprehensively assess the preparedness and response capabilities of existing military and civilian defense and public health enterprises or to determine how to address gaps lies outside the scope of this study; however, **exploration of the following areas is suggested to address some of the challenges posed by synthetic biology:**

- (a) **Developing capabilities to detect unusual ways in which a synthetic biology–enabled weapon may manifest.** For consequence management, expanding the development of epidemiological methods (e.g., surveillance and data collection) would strengthen the ability to detect unusual symptoms or aberrant patterns of disease. Enhancing epidemiological methods will have an additional benefit of strengthening the ability to respond to natural disease outbreaks.
- (b) **Harnessing computational approaches for mitigation.** The role of computational approaches for prevention, detection, control, and attribution will become more important with the increasing reliance of synthetic biology on computational design and computational infrastructure.
- (c) **Leveraging synthetic biology to advance detection, therapeutics, vaccines, and other medical countermeasures.** Taking advantage of beneficial applications of synthetic biology for countermeasure research and development is expected to prove valuable, along with corresponding efforts to facilitate the entire development process, including regulatory considerations.

Although addressing the potential concerns posed by synthetic biology in the age of biotechnology will remain a challenge for scientists and for the nation’s defense, there is reason for optimism that, with continued monitoring of biotechnology capabilities and strategic biodefense investments, the United States can foster fruitful scientific and technological advances while minimizing the likelihood that these same advances will be used for harm.

1

Introduction

Scientific advances over the past several decades have rapidly accelerated the ability to engineer existing living organisms and potentially create novel ones not found in nature. Synthetic biology collectively refers to concepts, approaches, and tools that enable the modification or creation of biological organisms. These concepts, approaches, and tools are being developed and refined by researchers in universities, governments, and industry in the United States and around the globe. Although synthetic biology is being pursued overwhelmingly for beneficial and legitimate purposes, such as addressing disease, remediating pollution, and increasing the yield of crops (see Box 1-1), there are potential uses that are detrimental to humans and other species. To inform investments to mitigate potential threats, those responsible for protecting the security of nations must consider how these emerging approaches and technologies might be used in acts of warfare or terrorism, the intent and capability of adversaries to effect such uses, and the potential impacts of such attacks.

Statements and reports issued over the past several years have come to different conclusions regarding the national security threats posed by emerging biotechnologies and the level of concern that is warranted. Former Director of National Intelligence James Clapper, in his 2016 annual threat assessment to Congress, grouped concerns about genome editing, an example of synthetic biology technology, under discussion of weapons of mass destruction (Clapper, 2016). Reports of federal government advisory committees, such as the 2016 report of the President's Council of Advisors on Science and Technology, "Action Needed to Protect Against Biological Attack" (PCAST, 2016), and a 2016 report of the JASON advisory group on potential implications of the gene editing platform CRISPR and other technologies for U.S. national security (Breaker, 2017), posit that biotechnology presents a new and significant threat. However, bioweapons are not a new phenomenon, and others have countered that, although advances in synthetic biology may add to the biological weapons landscape, these developments do not fundamentally change the landscape or warrant special action to address concerns (Vogel, 2013; Jefferson et al., 2014). That argument has been based on the notion that using natural pathogens to cause harm may be easier and just as effective as using synthetic biology to create bioweapons, and so synthetic biology did not change the level of concern, at least at that time (A. Paul interview with K. Vogel, February 24, 2006, New York, as cited in Vogel, 2012; Jefferson et al., 2014).

Although it is possible to imagine numerous types of malicious uses of synthetic biology, making informed decisions about whether and how to mitigate these potential uses requires a realistic assessment of the security concerns that this technology creates. To that end, the U.S. Department of Defense, working with other agencies involved in biodefense, asked the National Academies of Sciences, Engineering, and Medicine to develop

BOX 1-1 Benefits of Synthetic Biology

The field of synthetic biology opens tremendous possibilities for the application of biotechnology to improve human well-being, as well as the health of animals, plants, and the environment. Such applications hold substantial economic potential. For example, annual U.S. revenues from genetically engineered plants and microbes are estimated to exceed \$300 billion, and industrial biotechnology (the use of biological components to generate industrial products) is estimated to account for more than \$115 billion in annual U.S. revenues. New applications for biotechnology, particularly those driven by innovations in synthetic biology, are expected to further grow the size and reach of the bioeconomy (White House, 2012).

Often looked to as a means of producing products that would otherwise be difficult to obtain, synthetic biology has already led to new ways of producing pharmaceuticals including opioids and the antimalarial drug artemisinin. There are ongoing efforts to engineer microorganisms to produce fuels, act as detection devices, and clean up toxic spills. Synthetic biology is also seen as a potential means to grow organs for transplant, manipulate the microbiome, and even produce cosmetics. In addition to such application-driven goals, synthetic biology is also advancing the reach and role of science in society by inspiring more people to engage in biological experimentation, such as through the International Genetically Engineered Machine competition or by engaging with community laboratories. This broad array of applications and implications suggests that the potential benefits of synthetic biology are limited only by human creativity and imagination.

a framework to guide an assessment of the security concerns related to advances in the life sciences in the “age of synthetic biology,” to assess the level of concern warranted for various advances, identify areas of potential vulnerability, and provide ideas for options that could be considered to help mitigate potential vulnerabilities. To aid decision making in agencies across the biodefense enterprise, including the U.S. Department of Homeland Security, the U.S. Department of Health and Human Services’ Office of the Assistant Secretary for Preparedness and Response, the intelligence community, and other agencies, the Department of Defense asked the National Academies to consider potential concerns that are relevant to all U.S. citizens, both at home and abroad, in both civilian and military contexts. See Box 1-2 for the Statement of Task.

The study focuses on activities that could directly threaten human health or the capacity of military personnel to execute their missions. There are other conceivable uses of synthetic biology that are outside the scope of this study. The study does not address the potential ways in which plants, animals, and the pathogens that affect them could be modified for malicious purposes, for example, to undermine agricultural productivity, although the economic and societal impact of such an attack could be substantial. The study also does not address the modification of organisms to affect the environment or materials. Nonetheless, the technologies that might be used to threaten agricultural, environmental, or material targets, and the capabilities associated with those technologies, are likely comparable or even identical to the technologies and capabilities discussed in the report; as a result, the framework and analyses presented in the report may be useful for a broader array of contexts than those addressed in this study.

Finally, the report does not weigh the benefits on balance with the risks of synthetic biology advancements. Synthetic biology can play a role in achieving a number of societal goals but it is not within the purview of this study to compare the size or nature of those benefits with the potential risks. It is not the intent of the report or the study sponsor to imply that research efforts that use synthetic biology approaches for beneficial purposes should be curtailed.

BOX 1-2 Statement of Task

To assist the U.S. Department of Defense's Chemical and Biological Defense Program (CBDP), the National Academies of Sciences, Engineering, and Medicine will appoint an ad hoc committee to address the changing nature of the biodefense threat in the age of synthetic biology. Specifically, the focus of the study will be the manipulation of biological functions, systems, or microorganisms resulting in the production of disease-causing agents or toxins. The study will be conducted in two primary phases and will be followed by a workshop. Initially, the committee will develop a strategic framework to guide an assessment of the potential security vulnerabilities related to advances in biology and biotechnology, with a particular emphasis on synthetic biology.

The framework will focus on how to address the following three questions: What are the possible security concerns with regard to synthetic biology that are on the horizon? What are the time frames of development of these concerns? What are our options for mitigating these potential concerns? The committee will publish a brief interim, public report outlining the developed framework. This framework will not be a threat assessment, but rather, will focus on ways to identify scientific developments to enable opportunities that have the potential to mitigate threats posed by synthetic biology in the near, mid, and long term, with the specific time frames defined by the committee. The framework will lay out how best to consider the trajectory of scientific advances, identify potential areas of vulnerability, and provide ideas for potential mitigation opportunities to consider.

In Phase 2 of the study, the committee will use the outlined strategic framework to generate an assessment of potential vulnerabilities posed by synthetic biology. Inputs to this assessment may include information about the current threat, current program priorities and research, and an evaluation of the current landscape of science and technology. Conclusions and recommendations will include a list and description of potential vulnerabilities posed by synthetic biology.

UNDERSTANDING SYNTHETIC BIOLOGY

Biotechnology is a broad term encompassing the application of biological components or processes to advance human purposes, while synthetic biology is a narrower term referring to a set of concepts, approaches, and tools within biotechnology. A variety of perspectives has been offered to characterize the core principles of synthetic biology and the activities of its practitioners (see, e.g., Benner and Sismour, 2005; Endy, 2005; Dhar and Weiss, 2007), but there remains no universally agreed-upon definition (*Nature Biotechnology*, 2009). One distillation is that synthetic biology “aims to improve the process of genetic engineering” (Voigt, 2012). Chapter 2 provides additional detail on how synthetic biologists pursue that improvement.

A hallmark of synthetic biology is the use of concepts and approaches common to engineering disciplines. These can include standardization of components (e.g., well-characterized functions encoded by DNA), the use of software and computational modeling for designing biological systems from those components, and the construction of prototypes based on those designs. Synthetic biologists frequently apply such approaches in iterative Design-Build-Test cycles to accelerate progress.

This report takes a broad view of the field and does not attempt to narrowly define the term synthetic biology or to precisely separate it from other kinds of biotechnology. The concepts, approaches, and tools developed to advance synthetic biology will continue to be integrated more broadly into the life sciences toolkit and applied toward many biological research and biotechnology activities. Should a malicious actor seek to misuse such approaches, distinctions based on terminology will be irrelevant; similarly, the potential strategies for mitigating biodefense concerns are unlikely to be tied to a precise distinction between synthetic biology and other related activities. As a result, the analyses in the report focus on the potential applications of synthetic biology (also

described as synthetic biology–enabled capabilities or uses of synthetic biology) rather than on synthetic biology concepts, approaches, and tools themselves. In particular, the study was guided by the focus laid out in the Statement of Task on “the manipulation of biological functions, systems, or microorganisms resulting in the production of a disease-causing agent or toxin.” Modifying a pathogen to facilitate its rapid spread through a population, manipulating a biological system to produce a potent toxin, introducing antibiotic resistance into an infectious microorganism, and purposely weakening a person’s immune system are just a few examples of the potential types of malicious uses addressed.

ASSESSING POTENTIAL BIODEFENSE CONCERNS

A fundamental component of this study is to provide a basis for assessing potential areas of concern in the age of synthetic biology. Establishing a process for considering concern is important because it provides structure and transparency to the analysis of specific factors and how these factors contribute to an overall level of concern. It thus enables an assessment to more clearly convey the reasoning underlying judgments about potential concerns, increases consistency across assessments, and facilitates the comparison of assessments undertaken by different analysts or conducted at different times.

A number of possible approaches can be taken to develop such a process. The report presents a framework, which is largely a qualitative, multicriteria model, that could contribute to a qualitative, quantitative, or semi-quantitative assessment. As presented in Chapter 3, the methodology used to generate and apply this framework was informed by a review of existing frameworks, previous assessments, and related work relevant to biodefense, synthetic biology, and other biotechnology threats. Relevant documents include NRC (2004), IOM/NRC (2006), Tucker (2012), U.S. Government (2012, 2014), HHS (2013), Blue Ribbon Study Panel on Biodefense (2015), Royal Society (2015), Cummings and Kuzma (2017), and DiEuliis and Giordano (2017). Selected prior analyses are described briefly in Appendix B. The framework presented in the report was also informed by the expert judgment of committee members and input received during the course of the study.

The report also applies the proposed framework to analyze potential concerns associated with a number of synthetic biology–enabled capabilities. These analyses and their results are presented in Chapters 4–6. Detailed descriptions of how the framework was used to conduct the current assessment can help inform efforts to assess the significance of biotechnology developments that occur in the future; monitor key bottlenecks and barriers identified in the report that, if removed, could lead to a change in the relative level of concern; evaluate the change in the level of concern warranted when new experimental results are reported or new technologies arise; or scan the horizon to predict or prepare for potential future areas of concern.

While the report presents a framework for assessment of potential biodefense concerns and describes how that framework was applied to analyze synthetic biology–enabled capabilities, it is important to emphasize that this study is not a threat assessment. The study did not access intelligence or military information on potential actors, who may range from an individual to a dedicated team to a government body who may seek to misuse life sciences or their specific intent or specific capacity to undertake such misuse. Because information on actors is not included in the assessment presented in the report, a likelihood of harm cannot be fully estimated. By combining this assessment of concern with such classified information, however, the sponsor and others could, in the future, assess vulnerabilities and risks to inform decision making.

MITIGATING POTENTIAL BIODEFENSE CONCERNS

The report focuses on the state of science; it does not comprehensively assess the capability of the U.S. government to respond to the concerns identified in the report; it was outside of the study scope to access classified information or to comprehensively review the landscape of approaches being undertaken by the Department of Defense and other federal agencies to mitigate potential misuse of the life sciences. However, the existence and nature of anticipated mitigation options affects judgments about the levels of concern posed by synthetic biology capabilities. Thus, consideration of anticipated mitigation options is embedded in the framework presented in the

report, and the analyses presented include discussion of the potential for mitigating different synthetic biology-enabled capabilities based on an understanding of the current state of science.

The report also considers several types of mitigation approaches that may be useful for addressing some of the concerns arising from synthetic biology and biotechnology capabilities, as well as ways in which synthetic biology may affect those approaches (see Chapter 8). This portfolio of strategies includes options ranging from the promotion of norms of responsible conduct within the scientific community to strengthening the public health infrastructure to detect and respond to infectious disease outbreaks. However, because it was outside of the study's scope to consider all of the mitigation options available to the defense enterprise, the report does not make comprehensive, explicit recommendations regarding mitigation approaches.

STUDY APPROACH

To carry out the task, the National Academies appointed a committee including members with expertise in such areas as synthetic biology, microbiology, computational tool development and bioinformatics, biosafety, public health, and risk assessment (see Appendix D for biographical information).

The study was conducted in two phases. Phase 1 led to the development of an interim report proposing a framework for assessing potential vulnerabilities arising from developments in synthetic biology (National Academies of Sciences, Engineering, and Medicine, 2017a). The committee solicited feedback on the interim report from the synthetic biology, security, and policy communities to inform the second phase of the study. During Phase 2, the committee refined elements of the framework and applied the final framework to assess concerns posed by synthetic biology-enabled capabilities. This report, which represents the culmination of the study, presents the committee's assessment along with conclusions and recommendations. It thus extends and supersedes the interim report. This two-phase approach enabled the committee to understand the needs and motivations of the sponsor and other biodefense agencies, develop and refine a framework for assessing concerns, and apply the framework to provide an assessment of concerns associated with synthetic biology-enabled capabilities.

The study was informed not only by committee members' expert judgment, but also by the committee's analysis of information in published literature, including a review of existing frameworks and assessments as well as technical developments, progress, and barriers in synthetic biology, immunology, microbiology, and other relevant fields. The study was also informed by interactions with experts who shared their knowledge with the committee during public data-gathering meetings and webinars and by public comment and input. Additional details on the study process and data-gathering activities are provided in Appendix F.

The committee did not leverage classified information that others have created or utilized in their consideration of questions related to this study's task. Classified information was not included in the committee's deliberations; the resulting report is not classified and can be shared publicly. This facilitates the involvement of a wider community in the discussions during the study process and after the resulting reports are released. This report explores and envisions potential misuses of synthetic biology, including concepts that are regularly discussed in open meetings. The potential misuses as they are discussed in the report are neither comprehensive nor enabling in the level of information and detail provided; they are included to illustrate the expanding mission of biodefense in the age of synthetic biology.

Terminology

Although the report avoids precisely defining synthetic biology or drawing a strict distinction between synthetic biology and biotechnology, certain terms are used in a deliberate fashion to reflect the scope and nature of the assessment presented. For the purposes of this report:

- *Agent* or *bioagent* is used broadly to refer to any product created using biological components that may be intended to cause harm. In the context of synthetic biology, an agent could be a pathogen, a toxin, or even a biological component, such as a genetic construct or a biochemical pathway, that may be developed with the intent to harm a human target.

- *Actor* is used to refer to individuals or groups who may seek to effect an attack.
- *Target* is typically used to refer to the human beings harmed (or intended to be harmed) in an attack. In the context of manipulation of biological components, target may be used to refer to the intended outcomes of those manipulations.
- *Capability* is typically used to refer to the ability of an actor to produce and use an agent (or in some contexts, the ability for a target to mitigate adverse outcomes). The assessments presented in the report focus on synthetic biology-enabled capabilities, that is, applications that may be enabled by the misuse of synthetic biology concepts, approaches, or tools.
- *Vulnerability* refers to potential malicious capabilities against which we are not currently well protected. Vulnerabilities are a function of threat plus capabilities. Because the study did not include consideration of classified information about specific threats, specific actors, or specific capabilities within the U.S. government to address these threats, strictly speaking, it does not provide information on vulnerabilities but rather on *potential* vulnerabilities. Potential vulnerabilities are also referred to in the report as *concerns*.
- *Concern* is the term used to capture the committee's thinking regarding the defense implications of synthetic biology-enabled capabilities. *Level of concern* is used in reference to the relative intensity of the committee's opinion regarding potential misuse.
- *Threat* encompasses both an actor's capability to cause harm and the actor's intent to do so. Because the study did not include access to information on specific actors and their intent, the assessment produced is not a threat assessment per se. Rather, the report considers the types of malicious actions that could conceivably be taken and assesses the relative level of concern they pose.
- *Risk* refers to the likelihood and severity of harm. Again, because intelligence information on aspects such as actor intent was not considered, the likelihood of harm cannot be fully estimated and the term *risk* is not used in reference to the assessments undertaken as part of this study.

Organization of the Report

The report begins with a discussion of synthetic biology and explores how synthetic biology approaches are changing what can be accomplished by biotechnology (Chapter 2). The chapter highlights the fundamental Design-Build-Test cycle that characterizes a synthetic biology approach to problem solving. Appendix A discusses a number of concepts, approaches, and tools that are enabling continued progress in the field.

Chapter 3 describes the development of the framework presented in the report and provides information on the approach used in applying this framework to assess potential biodefense concerns posed by synthetic biology capabilities.

The following three chapters (4–6) discuss the results of the committee's assessment of synthetic biology-enabled capabilities including the use of pathogens as weapons (Chapter 4), the production of chemicals and biochemicals (Chapter 5), and the creation of bioweapons that alter the human host (Chapter 6).

Chapter 7 discusses advances in related fields whose convergence with synthetic biology may impact the ability to misuse biotechnology to create weapons, such as by helping to overcome challenges in delivery, stability, or targeting of an agent.

Chapter 8 discusses, from a broad perspective, some current approaches for mitigating concerns related to the malicious use of biotechnology, how synthetic biology may challenge those approaches, and conversely, how synthetic biology may help address challenges or bolster mitigation approaches.

Finally, Chapter 9 summarizes the relative concerns posed by the analyzed synthetic biology-enabled capabilities, highlights examples of key bottlenecks and barriers to monitor, and provides the report's conclusions and recommendations.

2

Biotechnology in the Age of Synthetic Biology

To frame and guide the study, the relationship of synthetic biology to other areas of biotechnology was explored along with the context in which synthetic biology tools and applications are being pursued. This chapter describes, in the context of this study, what it means to be in “the age of synthetic biology” and introduces key concepts, approaches, and tools that were considered.

WHAT IS SYNTHETIC BIOLOGY?

Biotechnology is a broad term encompassing the application of biological components or processes to advance human purposes. Although the term itself is thought to have been in use for only about a century, humans have used various forms of biotechnology for millennia. Synthetic biology refers to a set of concepts, approaches, and tools within biotechnology that enable the modification or creation of biological organisms. While there remains no universally agreed-upon definition of synthetic biology (with some defining it more narrowly and others more broadly; see, e.g., Benner and Sismour, 2005; Endy, 2005; Dhar and Weiss, 2007), one distillation is that synthetic biology “aims to improve the process of genetic engineering” (Voigt, 2012). By way of backdrop for this statement, it is useful to note that some of the concepts and approaches now associated with synthetic biology have roots going back to the early days of genetic engineering in the 1970s and the improvements and achievements that were envisaged then. In 1974, for example, the molecular biologist Walter Szybalski set the stage for some key synthetic biology concepts and presaged activities that have now been demonstrated.¹ An inflection point for the field occurred around the year 2000, after which synthetic biology gained significant attention and momentum. Two publications often identified with the field’s acceleration are by Elowitz and Leibler (2000) and Gardner et al. (2000). Although genetic engineering was occurring—and improving—prior to 2000, and the principles espoused by synthetic biologists were already noted and in use to varying extents (see, e.g., Toman et al., 1985; and Ptashne, 1986), that year marked a shift toward the adoption of approaches more typical of engineering disciplines, but which had previously been given only modest attention in the biological sciences.

¹ “Up to now we are working on the descriptive phase of molecular biology. . . . But the real challenge will start when we enter the synthetic biology phase of research in our field. We will then devise new control elements and add these new modules to the existing genomes or build up wholly new genomes. This would be a field with the unlimited expansion potential and hardly any limitations to building ‘new better control circuits’ and . . . finally other ‘synthetic’ organisms, like a ‘new better mouse’. . . . I am not concerned that we will run out [of] exciting and novel ideas . . . in the synthetic biology, in general” (Szybalski, 1974).

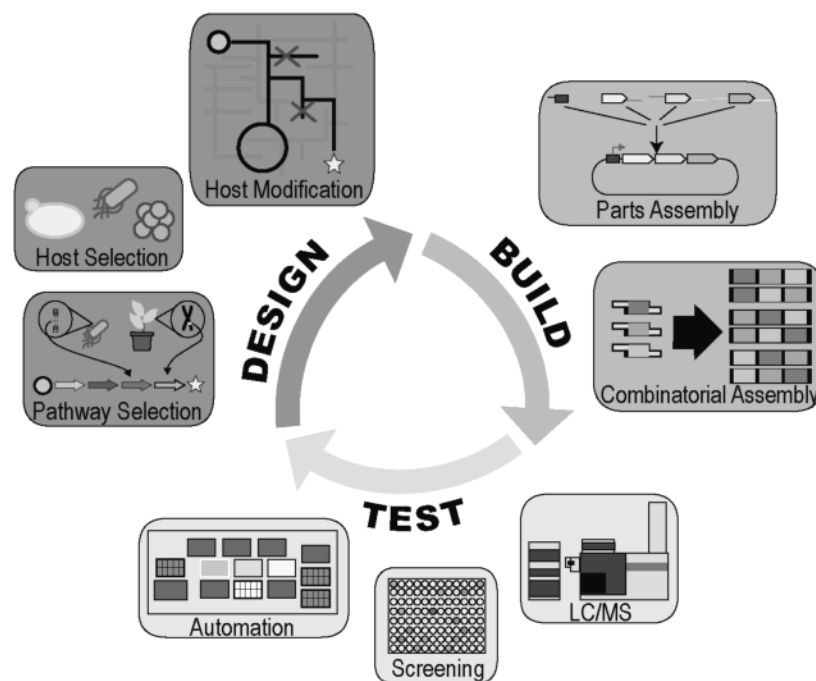


FIGURE 2-1 Design-Build-Test (DBT) cycle. This study approached synthetic biology concepts, approaches, and tools from the standpoint of their role in the DBT cycle, which is fundamental to synthetic biology.

NOTE: LC/MS = liquid chromatography–mass spectrometry.

SOURCE: Modified from Petzold et al., 2015.

In improving the process of genetic engineering, synthetic biology places special emphasis on the Design-Build-Test (DBT) cycle² (see Figures 2-1 and 2-2), the iterative process of designing a prototype, building a physical instantiation, testing the functionality of the design, learning from its flaws, and feeding that information back into the creation of a new, improved design. Developments such as enhanced computing power, laboratory automation, cost-effective DNA synthesis and sequencing technologies, and other powerful techniques to manipulate DNA have made it possible for biological engineers to rapidly repeat the DBT cycle to refine designs and products for a desired purpose. Key developments exemplifying these approaches include the establishment of standardized genetic parts registries, intensive use of models and other quantitative tools to simulate biological designs before building them, the availability of open-source DNA assembly methods, and the ability to create rationally designed genetic “circuits”—systems of DNA-encoded biological components designed to perform specific functions (Elowitz and Leibler, 2000; Gardner et al., 2000; Knight, 2003; iGEM, 2017a).

The age of synthetic biology is marked by the broad adoption and consolidation of these concepts, approaches, and tools within the DBT cycle to accelerate the engineering of living organisms. The concepts, approaches, and tools developed to advance synthetic biology will continue to be integrated more broadly into the life sciences toolkit and applied toward many biological research and biotechnology activities. As a result, this report does not draw a precise distinction between synthetic biology and other aspects of advancing biological sciences, but considers synthetic biology a crucial contributor to the spectrum of activities within biology and biotechnology more broadly.

² Sometimes referred to as a Specify-Design-Build-Test-Learn cycle or other variations.

The age of synthetic biology is ushering in not only novel technologies, but the application of engineering paradigms to biological contexts. The general intent to manipulate biological systems and to apply engineering paradigms from other disciplines is not new; from the introduction of recombinant DNA technologies in the 1970s to the present, there has been a concerted effort to manipulate genetic material and biological organisms. What has changed is the increased power of particular technologies that enable engineering paradigms to be applied to biological materials. Assessing new technologies and platforms that may enable the creative or destructive manipulation of biological materials, systems, and organisms will be important for identifying potential security opportunities and vulnerabilities.

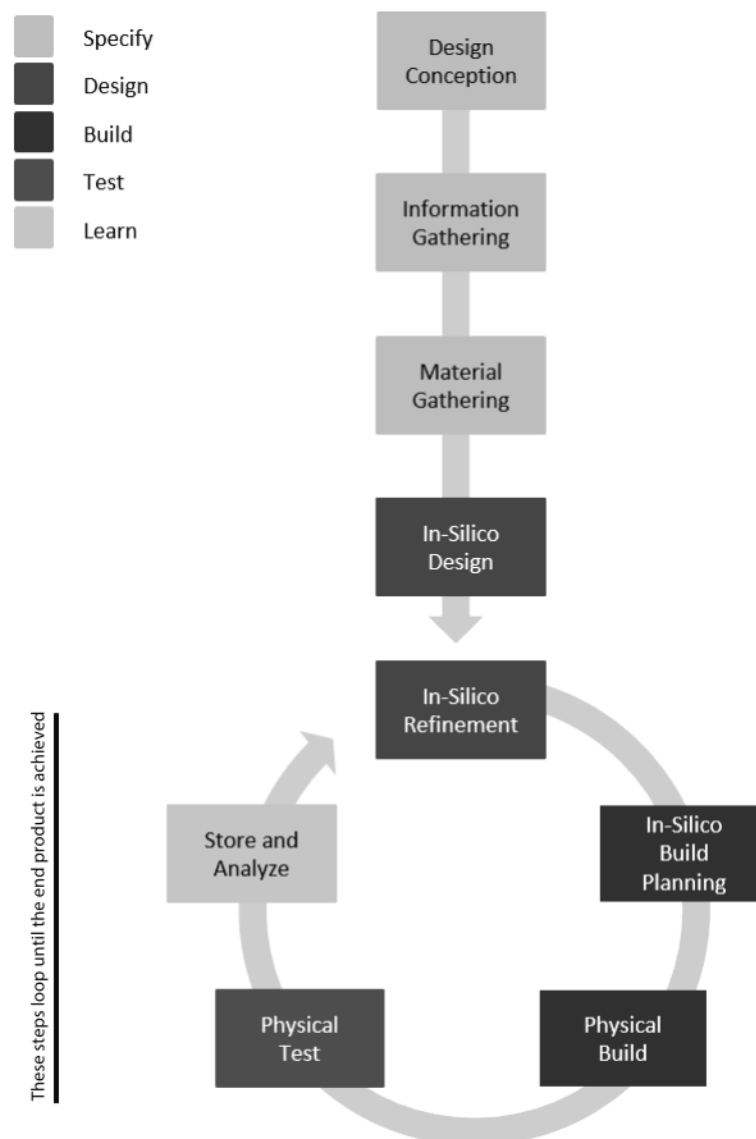


FIGURE 2-2 General workflow showing steps typical of the DBT cycle. This study focused on the core elements, Design-Build-Test, while recognizing that steps such as Specify and Learn can be considered separately or rolled into these core steps.

IMPLICATIONS OF THE AGE OF SYNTHETIC BIOLOGY

Synthetic biology is enabled by tools and techniques from a variety of scientific disciplines, from electrical engineering to computation to biology to chemistry. For example, the exponential improvements in DNA sequencing capabilities, initially developed to further our understanding of the human genome but soon applied to characterize many other organisms, have provided crucial raw material for synthetic biology and fueled innovation over the past decade. More recently, genome editing tools such as CRISPR/Cas9 (“clustered regularly interspaced short palindromic repeats”) (Jinek et al., 2012; Cong et al., 2013) have been adopted for synthetic biology techniques such as the regulation of gene circuits and the development of gene drives (genetic elements for which inheritance is favorably biased; see National Academies of Sciences, Engineering, and Medicine, 2016). Scientific progress in domains relevant to synthetic biology has been remarkably rapid; CRISPR/Cas9, for example, was extended from mammalian cell culture (in the United States) to primates (in China) in a single year (Cong et al., 2013; Jinek et al., 2013; Mali et al., 2013; Niu et al., 2014).

Two somewhat dichotomous phenomena are increasing the pace and progress of engineering of biological systems. The first is that bioengineering can be more theoretical, due to increased predictability of biological systems and evolving standards for biological performance. Biological engineering approaches make it possible to separate the design of a biological material or organism from its manufacture, and standards are evolving to facilitate a theoretical approach to biological design. Biological knowledge may thus be captured and applied in the design stage. The second phenomenon is the ability to try many different designs, often in parallel, and to potentially use directed evolution (see Appendix A) in living systems to perfect the design (see Box 2-1). The inexpensive technologies involved in designing and creating new DNA constructs to test make it easier to proceed without a hypothesis of how the design will work; in other words, it is “cheaper to make than to think.”³ However, the level of underlying biological knowledge still affects the degree to which these biological engineering techniques can be successfully applied; for example, adjusting well-understood pathways to increase ethanol production is fundamentally easier than increasing the virulence of *Francisella tularensis*, whose virulence mechanisms remain largely unknown.

These advances have real-world consequences for the development of new biotechnologies as well as their accessibility to actors of all types. On the positive side, it is expected that these technologies will enable a wider range of therapeutics, a wider range of biological detection and diagnostic methods, and opportunities to detect biological anomalies. However, these developments also potentially increase the power of even less-resourced malicious actors to produce a harmful biological agent. In this context, it is useful to consider the technologies that enable synthetic biology and how these developments may drive paradigm shifts in the practice of bioengineering.

Enabling Technologies for Synthetic Biology

Synthetic biology is enabled by numerous technologies that enhance success rates and facilitate experimentation, particularly in the DBT cycle. The development of these technologies to some extent defines the transition to the current age of synthetic biology. These include technologies specifically created for synthetic biology, as well as technologies developed for general molecular biology and biotechnology that are being exploited by synthetic biologists. These enabling technologies serve as the tools that facilitate the specification of biological designs and constructions. Key enabling technology areas, examples of which are described in more detail in Appendix A and below (see Specific Synthetic Biology Technologies and Applications), include the following:

- *DNA synthesis and assembly.* The heart of synthetic biology is the ability to make DNA constructs quickly and efficiently. Improvements in synthesis technology have followed a “Moore’s Law–like” curve for both

³ For example, researchers recently synthesized and tested more than 7,000 genes to identify diverse homologs capable of complementing the deletion of two essential *Escherichia coli* genes. While the function of those 7,000 genes could be inferred by sequence similarity, it was more tractable to prove their function via synthesis and testing rather than developing a model of their function from first principles. In practice, these large-scale efforts are synergistic with modeling techniques because they provide systematic data that can strengthen models for predicting biological functions (Plesa et al., 2018).

BOX 2-1 Designing Biology

Design in biology has traditionally differed from design in other engineering disciplines. In particular, biological design in the past has typically involved building and testing many designs to identify those that have the desired effect. The need for this trial-and-error process stemmed in part from the tools that were available; sequencing, synthesis, and gene editing tools have historically been too inexact and labor-intensive to permit systematic exploration of biological design spaces.

The complexity of biological systems makes it likely that biological design will continue to rely on trial and error, at least in part, for the foreseeable future. The balance between trial and error and explicit design is determined by our ability to predict phenotypic results from genotypic editing. Despite the continued need for trial and error, as the “craft” elements of genetic modification have been replaced with standards and practices, the discipline of design has come to play an increasingly key role in identifying strategies for specifying and building libraries that outperform previous approaches. In some cases, natural evolution can be co-opted to optimize designs by passaging samples through multiple generations of animal models or other living systems, where a selective pressure will identify the best constructs. In addition, aspects of biological systems can be discretely modeled with increasing accuracy. Examples of such advances include models of ribosome binding site strength (Salis Lab, 2017) and protein folding (Baker Lab, 2017), systems biology models (Palsson Lab, 2017), and statistical design tools (CIDAR Lab, 2016). None of these tools eliminate the need to build or test biological systems, but they reduce the size of the effective design space that must be explored to make progress toward a design goal. As tools supporting the building and testing of biological products improve in precision and throughput, larger design spaces can be explored.

The future of design in biology is expected to continue to separate the intent of the designer from the specification of genetic changes to make. Similar to the way that modern programming languages do not require software developers to understand how software routines are executed at the transistor level, biological design tools are becoming less dependent on base pair-level descriptions of genetic constructs. In other words, a synthetic biologist may not need to know the exact sequence of nucleic acids required in order to design a regulatory circuit for gene expression—simply specifying a particular goal, for example, the desire to integrate two predetermined biological signals, may be sufficient to return a blueprint for the Build stage. Importantly, design tools are not restricted to base pair-level descriptions of genetic constructs as output; they may instead output instructions for libraries of designs to build and test (e.g., suggesting a range of sequences to vary expression level of a regulatory protein) or conditions for mutagenesis, evolution, and selection (e.g., to augment rational design with directed evolution)—thus allowing the designer to more efficiently identify improved biological systems.

reductions in costs and increases in the length of constructs that are attainable. These trends are likely to continue.

- *Genome engineering.* Although in the past it has proven possible to engineer organismal or viral genomes via painstaking mutational methods, the ability to synthesize DNA quickly, coupled with improvements in transformation technologies and “booting” (the steps needed to go from DNA to a viable organism), has led to an acceleration in the ability to make mutations, including multiple mutations in parallel (e.g., Wang et al., 2009). In particular, the ongoing CRISPR revolution (Doudna and Charpentier, 2014) has led to the ability to introduce site-specific changes into a wide variety of organisms that may have previously been refractory to such techniques.
- *Improved computational modeling.* With new approaches to modeling biological systems and improved computing power, more complex biomolecular designs and system behaviors can be explored. This allows for larger areas of the theoretical “design space” in biology to be explored and tested in parallel, leading to better working systems in less time. Modeling advances are abetted by new computational advances

in machine learning and big data that have allowed the results of past experiments (both successes and failures) to inform the next round of design and experimentation. In the future, the creation of “rules” from the machine learning process should greatly improve the specification of future successful designs.

- *Genetic logic.* A key development in the field that meshes with improvements in modeling is the development of genetic logic circuits (Moon et al., 2012; Kotula et al., 2014) that allow living systems to make basic “decisions” based on both current inputs to the system (combinational logic) and the history of inputs (memory or sequential logic). The inherently programmable nature of genetic logic circuits is expected to mesh with advanced modeling approaches to improve the DBT cycle. An example of the use of genetic logic is plants that have been modified to act as radiation sensors capable of indicating when large amounts of gamma radiation have been detected (Peng et al., 2014).
- *Directed evolution.* While directed evolution methods are not new, their application has been accelerated by recent advances in DNA synthesis and genome engineering and are thus addressed in this report under the umbrella of biotechnology in the age of synthetic biology. Directed evolution methods stand both as an alternative to design-based models and as a supplement to them, in that they can return enormous amounts of data on fitness landscapes that can further improve computational modeling approaches. Additionally, the combination of design and selection moves constructs well beyond the bounds of what nature would attempt while still allowing the facile repair of unintended unnatural or less-fit deficiencies and interactions. A somewhat notorious example of the use of directed evolution was the introduction of an engineered version of a more virulent strain of influenza virus into ferrets, where it rapidly evolved to become airborne-transmissible (Fouchier, 2015). While this research was done for reasons some argue were appropriate, it also provided a blueprint for potential misuse.

Engineering Paradigms for Synthetic Biology

Enabling technologies have allowed synthetic biologists to make genetic changes in organisms with greater ease, precision, and scale. As a maturing engineering discipline, synthetic biology is also being advanced by engineering paradigms that allow these tools to be used with greater predictability of result. Engineering paradigms are methods of adapting enabling technologies to abstraction, standards, computing, workflow optimization, and other engineering principles. If enabling technologies provide options for *what* tools will be used in synthetic biology, engineering paradigms describe *how* these technologies will be used. In other words, these paradigms encompass the processes and decisions followed in designing, building, and testing biological constructs. The following engineering concepts and paradigms are particularly relevant to the context of this study:

- *Specify-Design-Build-Test-Learn cycle.* The Specify-Design-Build-Test-Learn cycle refers to an iterative process that requires a formal description of the desired biological behavior or function (Specify), the planned modification of an organism in silico or via rational design principles to realize that behavior (Design), the physical assembly of the biological material representing those designs (Build), the testing of the material to determine if it functions as specified (Test), and formally capturing and storing information about the entire process to inform the next revision or subsequent design (Learn). The boundaries between the cycle stages are fluid, and for the purposes of this report, the cycle is simplified to Design-Build-Test, with other stages implicitly included in these core elements. For example, Specify is incorporated into Design, and Learn is incorporated in the analytical steps of Test. Additional elements that are pertinent to biodefense considerations, such as Scale and Delivery, are also included.
- *Combinatorial approaches.* Although not an engineering paradigm per se, it is a fundamental shift that in many cases, it is now often “*cheaper to make than to think.*” It is becoming increasingly common to use combinatorial approaches—approaches in which a large number of genetic variants are created and then tested. Variants can be created by using a technique in which a large number of DNA variants are incorporated systematically to synthesize multiple variants (i.e., combinatorial assembly). The concept is that one can generate a large number of variants with limited knowledge of sequence-function relationships.

These approaches enable many design options to be explored, even in the absence of predictive tools to model the performance of those designs. Directed evolution is a related concept, discussed in Appendix A.

- *High-throughput data acquisition.* The speed of the DBT cycle has been greatly increased by the raft of enabling technologies such as combinatorial assembly (Smanski et al., 2014; Carbonell et al., 2016), CRISPR/Cas-based editing methods (Black et al., 2017; Schmidt and Platt, 2017; Mendoza and Trinh, 2018), and directed evolution (Cobb et al., 2013; Tizei et al., 2016). By synergizing with advances in analytical chemistry and biology, such as microfluidics and high-throughput sequencing, these technologies may allow the functional assessment of millions of constructs in parallel, hence providing particularly robust feedback for the next iteration of design.
- *Separation of design and manufacturing.* Specifying and designing a system can now be done in one location (e.g., an academic environment) while the manufacturing process (the Build step in the DBT cycle) is done in another location (e.g., a remotely operated facility or “cloud laboratory”). The increasing physical and virtual separation of design and manufacturing not only further increases the accessibility of synthetic biology but also creates potential security concerns where designs cannot necessarily be explicitly connected to manufacturing locations and vice versa.
- *Standards.* Standards have emerged that make DNA assembly easier and parts more “sharable” (e.g., Gibson and modular cloning assembly methods). Data standards such as Synthetic Biology Open Language⁴ have made the sharing, analysis, and software ecosystem of synthetic biology increasingly sophisticated. Such standards may ultimately allow engineers to focus on raising the level of abstraction in designs since lower-level mechanisms have been well defined and vetted.

SPECIFIC SYNTHETIC BIOLOGY TECHNOLOGIES AND APPLICATIONS

The technologies and engineering paradigms described above have led to a number of applications that drive synthetic biology development because they provide unique ways to take advantage of what synthetic biology offers. They are not all unique to synthetic biology, nor are they all routinely used to explore synthetic biology designs. For example, all synthetic biologists use software to store and analyze DNA sequences and use some form of computation in specifying designs (e.g., using biophysical models or algorithms to design ribosome binding sites, to check folding energies of DNA primers used for amplification and assembly, or to refactor the DNA sequence encoding a protein to increase protein production, a technique known as “codon optimization”). However, far fewer have the requisite library of DNA parts and accompanying software tools to achieve a level of abstraction that would allow the researcher to query, for example, a logic gate that accepts glucose concentration as input and activates transcription of a tethered reporter when a specific concentration is achieved. In other words, there are approaches and tools that are continuing to develop and gain traction within synthetic biology but which have not necessarily reached their full technical potential or user adoption.

Although the technologies used in each of the component phases of the DBT cycle may evolve over time or be replaced by new technologies, the fundamental concepts of the DBT cycle will stand. Thus, it is useful to consider current technologies and anticipated future developments in terms of the ways in which they enable the DBT cycle. However, it is important to recognize that the component phases of the DBT cycle are not strictly separate. It is possible, even probable, that some technologies or approaches will have impacts across multiple phases of the DBT cycle; one such example may be directed evolution, where repeated passage in a model host or in cell cultures under stress permits nature to Design, Build, and Test new phenotypes. There are also likely areas in which advances in synthetic biology capabilities relevant to biodefense would arise from synergies or convergence among technologies relevant to different phases. For example, it is important to consider potential synergies between Design technologies and Build technologies, because a malicious actor would need both Design and Build capabilities to carry out an attack. Similarly, synergies may arise if large-scale Test technologies are developed to match the enormous output of certain Build technologies, thus helping those Build technologies reach their full potential.

⁴ See <http://sbofstandard.org>. Accessed November 9, 2017.

TABLE 2-1 Synthetic Biology Concepts, Approaches, and Tools That Enable the DBT Cycle

| Key Synthetic Biology Concepts, Approaches, and Tools | Design | Build | Test |
|---|--------|-------|------|
| Automated biological design | | | |
| Metabolic engineering | | | |
| Phenotype engineering | | | |
| Horizontal transfer and transmissibility | | | |
| Xenobiology | | | |
| Human modulation | | | |
| DNA construction | | | |
| Editing of genes or genomes | | | |
| Library construction | | | |
| Bootstrapping of engineered constructs | | | |
| High-throughput screening | | | |
| Directed evolution | | | |

NOTE: Shading indicates which phase of the DBT cycle each example aligns with most closely. See Appendix A for full descriptions.

Appendix A describes a core set of current synthetic biology concepts, approaches, and tools that enable each step of the DBT cycle, focusing particularly on areas in which advances in biotechnology may raise the potential for malicious acts that were less feasible before the age of synthetic biology. Although the examples presented are intentionally quite broad and somewhat arbitrary—and do not represent an exhaustive list of all technologies or all possible applications of synthetic biology—they provide useful context for understanding how specific tools or approaches might enable the potential capabilities analyzed in Chapters 4–6 and can be adapted to assess new areas of concern as the biotechnology landscape continues to evolve. In addition, although Appendix A captures the main known technologies at the time of writing, this list will need to be updated and modified to stay relevant as the science advances.

Table 2-1 summarizes how the concepts, approaches, and tools described in Appendix A map to the phases of the DBT cycle. Going forward, it will be useful to consider how each phase of the DBT cycle may be further enabled by future developments in technology and knowledge, particularly in areas where a current bottleneck may be overcome. Appendix A also indicates the relative degree of maturity of specific techniques discussed (see Figure A-1).

3

Framework for Assessing Concern About Synthetic Biology Capabilities

The U.S. Department of Defense asked the National Academies of Sciences, Engineering, and Medicine to “develop a strategic framework to guide an assessment of potential security vulnerabilities related to advances in biology and biotechnology, with a particular emphasis on synthetic biology.” In public meetings, Department of Defense representatives clarified that the primary purpose of the framework was to serve as a tool to aid the consideration of the relative level of concern indicated for current and future synthetic biology-enabled capabilities. It was determined that the framework needed to be flexible enough to be applied in a variety of circumstances and for a variety of purposes, such as: analyzing existing capabilities to evaluate the level of concern indicated at present; understanding how various capabilities compare to, interact with, or complement each other in terms of their level of concern; identifying key bottlenecks and barriers that, if removed, could lead to a change in the relative level of concern; evaluating the change in the level of concern warranted when new experimental results are reported or new technologies arise; and horizon-scanning to predict or prepare for potential future areas of concern. This chapter describes the development of the framework and how it was used to facilitate an expert-based qualitative ranking of capabilities based on a well-defined set of factors to capture relative levels of concern.

APPROACH TO DEVELOPING THE FRAMEWORK

The process used to develop the framework generally followed best practices in expert elicitation and elicitation of attributes and value functions for multiattribute modeling (Morgan and Henrion, 1990; Clemen, 1991; Keeney, 1992; Keeney and Raiffa, 1993). First, the existing frameworks listed in Appendix B were reviewed, along with other published literature, to develop a list of factors that have been identified as being relevant to assessing concerns about the use of synthetic biology. A number of different frameworks have been developed to assess concerns associated with emerging technologies. In biology, these frameworks have typically assessed concerns based on features and capabilities of the biotechnology itself, particularly the capabilities the technology may provide to someone who would wish to create harmful biological entities for a specific malicious use. Some frameworks also consider the severity of potential adverse outcomes and the ability to manage them through detection, mitigation, or attribution. Other work has focused on assessing concerns associated with particular types of experimentation

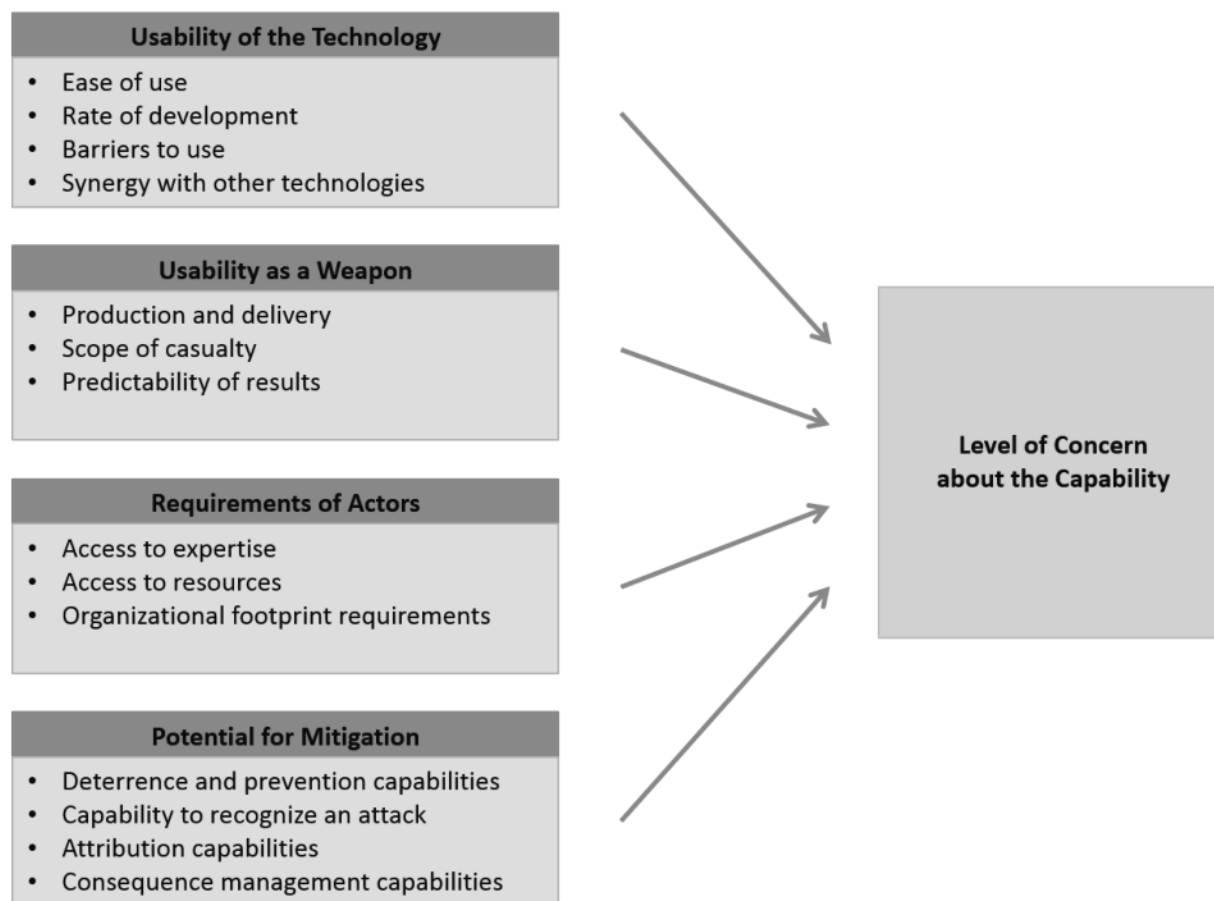


FIGURE 3-1 Framework for assessing concern. NOTE: The framework consists of four factors, along with descriptive elements within each factor, which delineate the information used to assess the level of concern for particular synthetic biology-enabled capabilities.

that may provide generalizable features applicable to a broader set of technological dual-use concerns.¹ Another framework approach, typically employed by security groups, is to use scenario-based assessments to identify potential vulnerabilities and the potential ways to mitigate them. Often referred to as “red-teaming,” this approach uses vignettes to describe details of a hypothetical scenario such as specific agents, actors, and affected populations. Although this approach can be informative, some scenario-based frameworks are hampered in the context of biodefense by a lack of evidentiary case studies and by the fact that one can come up with an almost limitless list of malicious activities that could potentially be pursued with biology (Lindler et al., 2005), and so the work is, by definition, never complete or comprehensive.

This review of the literature was followed by a process to identify terminology, factors, and approaches that resonated most within the context of the study charge. The outcomes of that process were formalized into a set of factors and elements within each factor, summarized in Figure 3-1 and described in more detail below.

¹ As defined by the National Science Advisory Board on Biosecurity, “Research yielding new technologies or information with the potential for both benevolent and malevolent applications is referred to as ‘dual use research.’” See <https://osp.od.nih.gov/biotechnology/nsabb-faq>. Accessed November 15, 2017.

These factors delineate the information that would be used to assess the level of concern for particular synthetic biology-enabled capabilities.

Developing quantitative or fixed scales for these factors was not attempted, nor was there an attempt to weight the factors relative to each other in terms of importance or impact on level of concern. Many of the factors and their descriptive elements are interdependent in that they capture ideas that are similar to or overlap with other factors and descriptive elements and are thus correlated with each other, requiring complex considerations for quantification. Instead, a qualitative approach was taken, using the factors and their descriptive elements to guide discussions and inform the assessment of relative level of concern for various synthetic biology capabilities. The assessment of each individual capability then fed into a holistic, relative ranking of the capabilities in terms of level of concern, similar to the methodology used in other studies (Morgan et al., 2001; Willis et al., 2004, 2010).

FACTORS FOR ASSESSING CONCERN

The framework for assessing concern consists of four factors, along with descriptive elements within each factor, as represented in Figure 3-1. The factors are usability of the technology, usability as a weapon, requirements of actors, and potential for mitigation. Conclusions about the relative level of concern about any particular synthetic biology capability are influenced by these four factors; in other words, capabilities that have lower technical barriers to use, more qualities that would enable use as a weapon, low actor requirements in terms of expertise or resources, and a low likelihood of mitigation would be of relatively more concern than capabilities for which there are high technical barriers to use, fewer qualities that would enable use as a weapon, high actor requirements in terms of expertise and resources, and a high likelihood of mitigation. As represented in this framework, those are the two extreme ends of the spectrum of concern. To complement and expand on the factors and descriptive elements, Appendix C lists illustrative questions that arose during the study process that can help facilitate the use of the framework.

Usability of the Technology

Biotechnology is a fast-moving field, and in some ways, synthetic biology is accelerating and broadening the usability of tools to achieve various capabilities. The first factor in the report's framework, usability of the technology, captures the idea that as tools become more usable, they become more accessible to more people, and therefore the concern about them being deployed for malicious use increases.

Four main elements were included in this study's assessment of the usability of technologies: ease of use, rate of development, barriers to use, and synergy with other technologies. Rather than attempting to formally score each of these elements for each capability analyzed, these elements were incorporated into one overall assessment of the usability of the technology for each capability considered.

Ease of Use

If a technology is easier to use, it is more likely to be used. Technologies that are in common use are likely to be more accessible and therefore more vulnerable to misuse, though it is also important to consider how outdated or less frequently used technologies may still be exploited for harm.

Advances in technology have made it easier to perform such tasks as creating single-nucleotide modifications and adding genes. Applications that employ combinatorial approaches to generate and test multiple design variants often involve complex work at large scales—as well as a high degree of unpredictability—thus putting them at the more difficult end of the spectrum. The availability of detailed information about a specific gene or pathway of interest also affects how easy or hard it is to use available technologies to manipulate that gene or pathway. These are the types of considerations that analysts can use to determine how much concern is warranted based on the ease of use of the technologies needed for a given application.

Rate of Development

All technologies follow some form of development curve over time. Technological capabilities that are developing rapidly are generally of more concern than those that are still far off in the future. If there is a known commercial use for a technology, private-sector investments may accelerate the rate of development, while technologies that do not have an identified commercial value may follow a slower path, advancing through smaller, disconnected efforts and public funding. Novel technologies may be characterized by rapid improvements in accuracy and throughput as their developers try to establish new markets or compete in existing ones. Technologies that have filled a unique market niche may survive for a long time with only minor improvements in scale or reductions in cost (e.g., the polymerase chain reaction, or PCR, has been in use for decades), while other technologies lose their prominence after being displaced by innovations (e.g., next-generation sequencing, also known as high-throughput sequencing, allows large numbers of genetic sequences to be determined far more rapidly than previous sequencing technologies and is expected to replace older technologies in some molecular identification applications).

Technologies for the synthesis of ever-larger DNA constructs are currently evolving rapidly, as are technologies for editing genes and genomes. For example, it is expected that the synthesis of all chromosomes from one strain of yeast is nearing completion. The engineering of plants to produce raw or finished chemical products is another area that is maturing rapidly. Assessing the degree to which the rate of development affects the level of concern warranted for a given use of technology should include consideration of both the pace of the technology's evolution and the speed with which it is being adopted.

Barriers to Use

It is also important to consider the presence of significant bottlenecks or barriers, which can lower the likelihood that a technology will be used. For example, key gaps in one aspect of the Design-Build-Test (DBT) cycle, such as Design knowledge, can significantly limit the potential for malicious use of a given technology and consequently lower the level of concern related to how that technology might be used in another phase of the DBT cycle, such as Build. Identifying barriers can also provide insight into potential rapid changes in what may be achievable once those barriers are overcome. This is an especially important consideration in areas of synthetic biology with strong drivers (e.g., beneficial uses attracting significant research) that are pushing the barriers to be broken. Major technological leaps have the potential to change synthetic biology quickly and open up new possibilities; for example, Gibson Assembly® (Gibson et al., 2009) led to a sea change in the ability to compile genetic fragments.

Synergy with Other Technologies

Some technologies may be substantially enhanced by synergies with other technologies, leading to higher level of concern for the capabilities they may enable. For example, CRISPR/Cas9 can be used alone to make a specific modification to a targeted gene. But when CRISPR/Cas9 is coupled with emerging technologies for single-cell sequencing, it is possible to create random libraries of CRISPR/Cas9 guide RNAs, apply them in parallel to single cells, subject the cells to environmental pressures, and use single-cell next-generation sequencing to identify the “winners” (Datlinger et al., 2017)—a far more complex proposition than could be achieved with CRISPR/Cas9 alone.

In the field of computing, the semiconductor technology evolution has brought ever-greater computing power and data storage at ever-lower costs. At the same time, the evolution of networking technology has converged with computing to make computing more ubiquitous, powerful, and inexpensive, thanks in part to a concerted effort to identify and overcome bottlenecks and barriers in both computing and networking. Synthetic biology and sequencing technology may well show a similar convergence in the coming years, in which advances in annotation and predictable sequence-structure-function relationships lead to the ability to reliably design increasingly complex biological systems (Brophy and Voigt, 2014; Chao et al., 2015).

Such developments would have implications for both beneficial and malicious uses of synthetic biology

technology. In determining the level of concern warranted for any given capability, it is useful to consider how synergies among relevant technologies may create opportunities for new types of applications in the future. It is also useful to consider how a breakthrough relevant to one aspect of the DBT cycle might synergize with technologies relevant to other aspects to enable applications that were not previously achievable.

Usability as a Weapon

A central question is whether a capability enabled by synthetic biology can be used in such a way as to cause harm—that is, whether a capability can be used as a weapon. A great deal of previous work has sought to characterize what makes a substance “weaponizable” (Kadlec and Zelicoff, 2000; U.S. Congress, 2006; Carus, 2017). Drawing on that work, usability as a weapon was identified as a primary factor in the framework for assessing concerns related to synthetic biology-enabled capabilities. A capability determined to have more characteristics that make it usable in the development of a weapon warrants a higher level of concern than a capability with fewer characteristics for that purpose. In particular, the elements considered as part of usability as a weapon include implications for production and delivery of a weapon, the expected scope of casualty for a given use of technology, and the predictability of the intended results.

Production and Delivery

There are two types of questions to consider with regard to the production and delivery of weapons created with synthetic biology. They build upon a large body of existing work related to the classical understanding of the use of pathogens to create weapons of mass destruction. Previous frameworks for understanding threats related to bioweapons outline a series of key steps involved in creating a bioweapon and using it in an attack. These steps include bioagent production, stabilization, testing, and delivery (van Courtland Moon, 2006) and might include specific processes such as growing large amounts of an agent, milling it into a powder form, making the agent stable enough to be sprayed in a crop duster or withstand other means of mass dispersal, and testing its effectiveness in animal studies. These steps were considered significant barriers to the production of bioweapons in the Cold War era, in effect limiting bioweapons capabilities to a few well-resourced nation-states. In assessing the biodefense concerns posed by biotechnology, it is important to consider (1) whether synthetic biology could lower the barriers related to bioagent production, stabilization, testing, and delivery or (2) whether advances in biotechnology areas other than synthetic biology may impact the potential to weaponize products created with synthetic biology.

The first item has to do with whether synthetic biology makes unnecessary any of the classically defined steps to weaponization and thus eliminates barriers previously associated with that step. For example, synthetic biology could potentially be used to enhance existing pathogens or create new ones, but it also raises the possibility of types of attacks in which the “weapon” involved is not a pathogen per se, but a genetic construct, toxin, or other entity. Deploying such alternative bioagents might not require the same type of large-scale production or purity of pathogens required for some traditional bioweapons. In addition, synthetic biology could raise concerns about smaller types of attacks that do not require mass dispersal, which could change the equation with regard to the need for stabilization. All of these elements could potentially reduce or eliminate barriers that previously were thought to hinder the use of bioweapons, so their presence would generally increase the level of concern.

The second item relates to how advances in other areas may impact the potential to weaponize products created with synthetic biology. For example, it may be important to consider how advances in technologies such as bioreactors² may change the nature of the production facilities required to produce harmful agents using synthetic biology.

² Bioreactors are vessels in which biologically active substances produce substances or biological components, a type of biotechnology that is not exclusive to synthetic biology.

Scope of Casualty

The scope of casualty it is possible to generate by using a synthetic biology capability to create a weapon gives a sense of the scale of the potential threat it poses. For capabilities that could lead to a large number of people impacted and/or a severe outcome like permanent disability or death, the concern level would be higher.

Predictability of Results

Predictability of results describes the degree to which a malicious actor could be confident that the intended result will be achieved when using a given technology to develop a weapon. A higher degree of predictability would be associated with a higher level of concern. While some technologies, applications, and types of attack may require extensive testing in order to ensure the intended impact, there may be a lower barrier to success if, for example, the bioagent would only need to be produced one time to have the desired outcome, if the attacker has the opportunity to deliver the agent multiple times, or if the attacker can create many versions of the agent to maximize the likelihood of success. To assess the overall predictability of results for the malicious use of synthetic biology, it is useful to consider both a need for testing and phenotype predictability.

Testing A large-scale, long-term, and highly resourced bioweapons operation could likely be expected to perform testing prior to deployment to ensure that the scaled-up bioagent behaves as intended and that the delivery or dissemination method is functional. This process would typically involve testing in animal models to ensure illness or lethality, as well as field testing in specific environments to ensure that the agent survives well enough to persist and infect targets. In the context of a synthetic biology-enabled weapon, it is useful to consider the degree to which testing would be necessary for a given use and how this testing might be carried out. If significant testing is not likely to be necessary, the concern would be higher.

Phenotype Predictability A related question is whether the genotype of a bioagent could be predictably engineered to yield the desired phenotypes. For example, are there known engineering strategies or preexisting research that outlines methods to predictably produce the desired result? Or can the properties of a bioagent be modeled with computational tools? The ability to predictably design, model, or construct an agent could reduce the need for testing. Agents with predictable genotype-phenotype relationships may also require fewer resources to deploy, since it may not be necessary to test multiple genotypes to obtain the desired phenotype. Therefore, as phenotype predictability increases, so does the level of concern.

Requirements of Actors

Any discussion of the concerns related to the potential malicious use of a specific biotechnology needs to include consideration of requirements of the person or people who would be involved in perpetrating an attack, here referred to as actors. Actors may range from a single individual to a dedicated team to a government body. They may be amateurs, biotechnology experts, or engineers or have some other type of relevant expertise. The complexity involved in exploiting a technology (see Usability of the Technology, above) will have varying impacts on the likelihood of use and therefore on the level of concern, depending on the capabilities of the actors. For example, whereas it may be impractical (or would take an extremely long time) for an individual actor to gain the necessary capabilities and knowledge to use a given capability to cause harm, a dedicated team might have the diversity of expertise necessary to enact the same plot much more quickly.

When analyzing how the requirements of actors affect the level of concern about a given capability, it is useful to consider questions related to the expertise an actor would need to possess to effect a given attack, the accessibility of the required resources, and the organizational footprint and infrastructure that would be required.

In addition, while this study did not include consideration of the intents or actual capabilities of actors, which would likely have required access to classified information, such information could, in the future, be incorporated into an assessment of vulnerabilities to inform decision making.

Access to Expertise

Some types of applications of biotechnology require a great deal of expertise in one or more areas, while other uses may require less expertise. The degree to which expertise requirements represent a barrier to malicious use of a technology depends on the expertise possessed (or obtainable) by a malicious actor. It is important to assess the gap between the types of expertise required and the types of expertise that actors might be expected to have access to. In some cases, exploiting synthetic biology for harm may require an actor to interact with the conventional research community to acquire goods, services, or expertise, in which case the concern would be lower because this would be a barrier that may enable malicious use to be detected earlier.

Access to Resources

The particular resources needed to effect a given malicious use of synthetic biology depend on many factors. Resource requirements can include money, time, laboratory equipment and other infrastructure, reagents and other raw ingredients, personnel and expertise, and other types of resources. If more resources are needed, the concern level is decreased because that reduces the number of potential actors. If fewer resources are needed, then there is a higher level of concern.

There are multiple, hypothetical ways for an actor to obtain resources. For example, if an actor requires the use of an expensive DNA synthesizer but lacks sufficient funds to purchase a new instrument via conventional channels (or fears an outright purchase would lead to discovery), the actor may consider purchasing a used synthesizer, obtaining legitimate or covert access to equipment at a company or university, coercing an innocent person with legitimate access to perform the work (via bribing, subversion, blackmail, or threats of harm), or resorting to outright theft. A solo actor could be better funded than a group sponsored by a poor nation-state. Conversely, a poor but resourceful actor might find ways to access even highly sophisticated technologies, for example, by enrolling in a graduate degree program, getting a job in a biotechnology company, or taking advantage of relevant service providers or brokers of services. Assessing needed access to resources is not always a straightforward proposition, but it is nonetheless an important consideration when evaluating potential concerns.

Organizational Footprint Requirements

If achieving a particular malicious use of synthetic biology requires a large organizational footprint, the concern will be lower compared to capabilities for which only a small organizational footprint is needed. Some malicious uses of synthetic biology might be achievable by an individual working with basic supplies and a rudimentary laboratory, whereas other types of attacks might require a larger organization, more personnel, or more extensive infrastructure. Furthermore, considering the organizational footprint that would be required to effect a given type of attack can shed light on the relative importance of other actor attributes, such as access to resources. Organizational footprint also affects considerations related to the potential for mitigation, such as the ability to identify suspicious activity and prevent an attack or the ability to attribute an attack to the actor responsible (discussed further under Capability to Recognize an Attack and under Attribution Capabilities, below). For example, activities requiring less equipment may be able to be pursued by actors with fewer resources and may be conducted in a clandestine laboratory, making detection or attribution more difficult and therefore making concern higher. Malicious uses requiring a large organizational footprint, on the other hand, might require an actor to have access to more funding or access to legitimate infrastructure (such as by being embedded within a university laboratory), increasing the likelihood of detection or attribution and leading to a lower level of concern.

Potential for Mitigation

The impact of an attack depends both on the actor's ability to deploy a weapon and on the target's ability to prevent, detect, respond to, or withstand the attack. To comprehensively assess concerns, it is important to consider mitigating factors that may diminish the likelihood that a synthetic biology capability will be effectively used to cause harm or that may reduce the damage caused. Elements within this factor include the ability to deter or prevent an attack, the ability to recognize when an attack has occurred, the ability to trace an attack to the responsible actor (or "attribute" an attack), and the ability to manage the consequences of an attack. Because this factor is a core part of the framework, considerations related to the potential for mitigation were included in the assessments of specific capabilities presented in Chapters 4–6; however, significant data gathering on U.S. mitigation capabilities was outside of the study scope and the assessments presented in those chapters are intended to be illustrative and to demonstrate the assessment process rather than provide a full analysis. Mitigation capabilities are also discussed further in Chapter 8.

Deterrence and Prevention Capabilities

Various factors can affect the likelihood that a malicious actor will decide to pursue an attack and then successfully execute it. One important element that is understood to deter adversaries from pursuing some types of biological attacks is the availability of countermeasures that limit the amount of harm an attack would cause. For example, the fact that the United States has smallpox vaccine stockpiled—and would thus have a ready countermeasure against an attack using smallpox—is expected to deter malicious actors from perpetrating attacks using smallpox.

One approach that has been used as a preventive measure is the establishment of regulatory and statutory safeguards that limit the ability to access particular pathogens or technologies and use them for harm. For example, by limiting access to certain pathogens, the Federal Select Agent Program is intended to reduce the likelihood of those pathogens falling into the hands of malicious actors who might seek to use them as a weapon.

In addition, activities such as intelligence gathering can contribute to deterrence and prevention by increasing the capacity to identify suspicious activities and intervene before an attack takes place, or to catch and punish an actor after an attack has occurred, as discussed under Capability to Recognize an Attack and under Attribution Capabilities, below. Intelligence gathering allows authorities to recognize and respond to activities that may indicate that an actor is preparing for a biological attack, such as by monitoring individuals or groups with a known intention to carry out an attack, monitoring individuals or groups with access to equipment or expertise necessary to develop a bioweapon, or tracing the procurement of supplies that could be used in a biological attack. However, because biotechnology is used for so many beneficial applications and because different combinations of technologies can be used for the same or different purposes, it can be challenging to identify activities, specialized equipment, or other signatures that distinguish suspicious activity from benign activity.

Capability to Recognize an Attack

In general, there is a higher level of concern about attacks that would require some time and work to identify (as a health threat and/or as a purposeful attack) compared with attacks that would be readily recognizable. Once an attack has occurred, recognizing the emergence of an unusual cluster of disease is the first crucial step toward launching an effective response. In addition, being able to differentiate between a natural disease outbreak and purposeful use of a bioagent is vital to preventing subsequent attacks and finding the perpetrators. This knowledge also can inform how medical personnel, public health organizations, and law enforcement or military authorities act to contain the scope of the damage. Public health programs and disease surveillance systems such as those under the purview of the U.S. Centers for Disease Control and Prevention are designed to facilitate the rapid identification and characterization of known infectious disease threats as they emerge. It is important to consider how synthetic biology might affect the ability to identify suspicious activity, recognize when an attack has occurred, and identify the individuals or groups that have been targeted.

Attribution Capabilities

The ability to attribute an attack to the actors responsible is crucial to consider as part of the framework, because attribution may provide a disincentive to attacks in some circumstances. That is, actors may choose different courses of action if their actions could lead to prosecution or retaliation; thus, there is a higher level of concern about attacks that would be more difficult to attribute. Attribution considers scientific evidence, its validation, and nonscientific types of information. In the future, it may be important to consider how attacks that use synthetic biology approaches could conceivably be amenable to the development and validation of different lines of molecular evidence. Such potential opportunities are discussed in Chapter 8, such as next-generation DNA sequencing and analysis of “scars” left by engineering techniques (e.g., a remnant of a DNA vector used to insert synthetically derived biological components).

Consequence Management Capabilities

Protocols and procedures for responding to public health emergencies and to biological and chemical attacks exist in both the civilian and military arenas (CDC, 2001, 2017d). These procedures often involve, for example, epidemiological methods of identifying victims, agents, and modes of transmission, as well as activities such as the development and use of vaccines, drugs, and antitoxins to save lives. Other relevant capabilities include emergency response capacity, availability of supportive healthcare facilities, and effective procedures for isolation and quarantine. When assessing the level of concern about any particular capability, it is important to understand how that capability could change the ability to mitigate the negative impact of an attack.

APPLYING THE FRAMEWORK IN THE ASSESSMENT OF CONCERN

The framework was developed both to facilitate the analysis of synthetic biology-enabled capabilities presented in subsequent chapters of this report, as well as to aid others in their consideration of current and future synthetic biology capabilities. To support and inform the application of the framework by other parties, this section describes the approach taken to identify potential areas of concern, the steps used to apply the framework, and key considerations that guided the analysis.

Approach Taken to Identify Potential Areas of Concern

A number of technologies support various aspects of the synthetic biology Design-Build-Test cycle; selected examples are captured in Appendix A. The interim report (National Academies of Sciences, Engineering, and Medicine, 2017a) released as part of this study identified these technologies as potential items for which the framework could be used to assess concern. However, the technologies themselves pose no inherent harm, and it would generally take a collection of technologies to create a specific capability that warrants concern. As a result, this final report describes how the framework was applied to assess *capabilities* (rather than *technologies*) that potentially pose a concern because of the harm they might enable.

A list of potential capabilities to evaluate was identified by gathering a range of possibilities that have been mentioned in various venues as potential concerns associated with synthetic biology and augmenting that list with additional possibilities that had not been previously raised. These potential capabilities were grouped into categories to ensure a consistent approach to their evaluation using the framework. The following potential capabilities were analyzed (see Chapters 4–6):

- *Re-creating known pathogenic viruses*: Constructing a known, naturally occurring pathogenic virus from the starting point of information about its genetic sequence.
- *Re-creating known pathogenic bacteria*: Constructing a known, naturally occurring pathogenic bacterium from the starting point of information about its genetic sequence.
- *Making existing viruses more dangerous*: Creating a modified version of a known virus in which one or more traits have been altered to make the virus more dangerous (such as by enhancing its virulence).

- *Making existing bacteria more dangerous:* Creating a modified version of a known bacterium in which one or more traits have been altered to make the bacterium more dangerous.
- *Creating new pathogens:* Constructing a pathogen from the novel combination of multiple parts, which may be derived from various organisms, designed computationally, or created through other strategies.
- *Manufacturing chemicals or biochemicals by exploiting natural metabolic pathways:* Producing a naturally occurring product, such as a toxin,³ by engineering an organism (e.g., bacterium, yeast, or alga) to contain the known biosynthetic or metabolic pathway for the desired product.
- *Manufacturing chemicals or biochemicals by creating novel metabolic pathways:* Creating a new biosynthetic pathway that enables an engineered organism to produce a chemical that is not normally produced biologically.
- *Making biochemicals via in situ synthesis:* Engineering an organism, such as a microorganism that can survive in the human gut, to produce a desired biochemical and delivering this microorganism in such a way that it can produce and release this product in situ.
- *Modifying the human microbiome:* Manipulating microorganisms that form part of the population living on and within humans—for example, to perturb normal microbiome functions or for other purposes.
- *Modifying the human immune system:* Manipulating aspects of the human immune system, for example, to upregulate or downregulate how the immune system responds to a particular pathogen or to stimulate autoimmunity.
- *Modifying the human genome:* Creating changes to the human genome through addition, deletion, or modification of genes or through epigenetic changes that modify gene expression. A subset of this category is the modification of the human genome through *human gene drives*, the incorporation of certain types of genetic elements into the human genome that are designed to pass from parent to child during reproduction and that would spread a genetic change through the population over time.

Steps Used to Apply the Framework

The framework is designed to facilitate a thorough analysis of any particular capability by providing a set of key factors to consider and specific elements to consider for each factor. To inform decisions, however, it is useful to consider capabilities in relation to each other, that is, to assess areas of concern in relation to other potential concerns. To that end, the framework was applied using the following steps, which can be followed by other, future framework users:

1. Gather and organize information about a capability in terms of the four framework factors and the elements relevant to each factor.
2. Compare information about the capability to information about other capabilities to determine how the level of concern for a given capability compares to the level of concern for other capabilities.
3. Consider all capabilities holistically, using the framework to inform judgments about relative levels of concern, based on all the information generated in steps 1 and 2.

Different types and levels of expertise may be required to successfully analyze the factors and elements related to any particular capability. This committee benefited from a wide range of expertise areas, including synthetic biology, microbiology, computational tool development, bioinformatics, biosafety, public health, and risk assessment.

For the first step, a qualitative approach was used to “score” each capability on each factor using a relative scale from low to high. For example, for the factor usability of the technology, the scale ranged from relatively low usability (which corresponds to relatively lower concern because it is relatively more difficult to use) to relatively high usability (which would be of relatively higher concern because it is relatively less difficult to use).

Figure 3-2 shows the first step in the process using an illustrative example. For the first capability, “Capability 1,” information associated with the elements relevant to the first factor, usability of the technology (which

³The phrase “chemical or biochemical” throughout the report includes toxins.



FIGURE 3-2 Capability 1 assessed with regard to usability of the technology.

includes ease of use, rate of development, barriers to use, and synergies with other technologies) was discussed and analyzed. Using that information, Capability 1 was placed on a relative scale ranging from low to high usability. Capability 1, the first capability discussed, was placed near the middle of the scale.

Next, another capability, “Capability 2,” was placed on the scale. To do this, each of the elements for the usability of the technology factor were discussed for Capability 2 and compared to those elements for Capability 1. A facilitated discussion was used to place Capability 2 on the scale relative to Capability 1 (see Figure 3-3). Note that the bar for Capability 2 is wider than the bar for Capability 1 in order to represent a broader range of concern regarding usability of the technology for Capability 2.

Each capability was considered in turn, with available information on each of the elements carefully discussed, reviewed, and compared to the corresponding elements for other capabilities, to place the remaining capabilities on the scale, as shown in Figure 3-4.

This process was repeated for each capability and each factor (Usability of the Technology, Usability as a Weapon, Requirements of Actors, and Potential for Mitigation). As the work progressed, the definitions of some of the factors and capabilities were refined, and adjustments were made to the assessments based on those refinements.

To help translate these graphics into usable information, five categories were created along the x axis: high, medium-high, medium, medium-low, and low. These categories are intended to reflect relative levels of concern,



FIGURE 3-3 Capability 1 and Capability 2 assessed with regard to usability of the technology.



FIGURE 3-4 All capabilities assessed with regard to usability of the technology.

not absolute levels of concern. No numerical scores were assigned to these categories and there was no attempt to normalize categories across factors (that is, to ensure that “medium” on one factor meant the same thing as “medium” on another factor) because such steps were not necessary for their use. Rather, capabilities were placed in the same category when they were seen as similar with regard to that factor. Not requiring the categories to have numerical meaning made it more straightforward to achieve agreement among the experts on the committee, with no loss of value in the information generated since all of the judgments were relative.

As a final step, all of this information was integrated into a holistic assessment of the relative levels of concern across the full landscape of capabilities considered. Chapter 9 presents the results of this holistic assessment (see Figure 9-1).

Key Considerations That Guided the Assessment

As described above, an expert-driven, qualitative, multiattribute methodology was used to develop the framework and apply it to assess concerns associated with synthetic biology capabilities. There are strengths and weaknesses of any methodology. The following considerations guided the assessments presented in this report and could help inform future users of the framework:

1. *The factors were consistently applied.* Care was taken to ensure that the factors were consistently used and appropriately incorporated into an assessment of overall level of concern. Each factor was reviewed separately for each capability and the entire list of capabilities was reviewed as part of the process of determining where each one belonged on the relative scale from “lowest concern for this factor” to “highest concern for this factor.” These graphs did not have absolute values but were maintained in relative terms, so that each capability was assessed relative to the others with regard to each factor. This approach reflects the level of precision that was included in the deliberations about the capabilities.
2. *The final assessment incorporates a holistic evaluation.* A holistic consideration of relative concern is a critical part of ensuring that the final ranking captures the full extent of the input from the ranking process. The relative placement of each capability on the scale of each factor is not deterministic of the final ranking, but rather provides consistent information to be used in making holistic judgments. The final rankings cannot be calculated based solely on the individual factor rankings since additional information may be brought to bear on that holistic judgment; the factors included in the framework are meant to inform holistic judgment, not to replace it or provide a checklist approach. However, the holistic assessment was grounded by consistent use of the factors; to maintain robustness of the factors, when a capability was placed on the scale of overall concern, it was compared to the ratings of the other capabilities already placed on the overall concern graph. For example, if Capability 1 was scored as a medium level of concern with regard

to usability of the technology and Capability 2 was scored as a relatively high level of concern with regard to usability of the technology, this information informed the assessment of overall level of concern about Capability 2 relative to Capability 1.

3. *The factor scaling approach has implications for future comparative assessments.* The factors that make up the framework were constructed specifically for this study and were refined through the process of applying them to assess specific capabilities. Using a relative scaling approach allowed these definitions to be refined and aligned as the study progressed. In addition, the use of a relative rather than absolute scale for the factors means that the placement of capabilities already on the scale may need to be adjusted as subsequent capabilities are assessed. For example, if a capability is introduced that holds a much greater concern than the highest-ranked item already assessed, either the already-assessed item might need to be moved down the scale or the scale might need to be extended to allow the new capability to be ranked as “very high” concern. An alternative approach that could be used in future assessments, rather than starting with any capability and making all subsequent judgments relative to it, could be to identify the highest and lowest capabilities on each factor, assign the highest “100” and the lowest “0,” and place all other capabilities on the scale relative to those capabilities.
4. *Choices may need to be made to capture uncertainty and variability.* In placing synthetic biology capabilities on low-to-high scales for each framework factor, placement reflected the range of potential concerns for a given capability, with particular exceptions noted in the analyses presented in Chapters 4–6. Uncertainty and variability beyond notable exceptions were captured by varying the width of the bar (see Figures 3-2 to 3-4 for notional examples), with a wider bar representing greater uncertainty or variability. During the assessment process, one case (re-creating known pathogens) initially had a very wide bar when assessing some of the factors, primarily because of the diversity of organisms that the capability included. In response, that capability was divided into two capabilities that were assessed separately (re-creating known pathogenic viruses and re-creating known pathogenic bacteria) to allow the assessment to be more precise.
5. *A qualitative assessment approach was used; other approaches to using the framework are possible.* Methodologies for technical forecasting in emerging areas such as synthetic biology are evolving to meet the needs of decision makers. The report uses the framework to conduct a qualitative assessment; other users could choose to apply the framework in different but still meaningful ways. In the future, other users may decide to pursue a more quantitative approach to conducting the assessment or to extend the framework to incorporate sources of information outside the study’s scope (such as intelligence on actor intent or additional information on U.S. mitigation capabilities). The choice to use a qualitative or quantitative approach would be impacted by the amount and types of information available and the level of precision and understanding that would be consistent with the available information. Were a quantitative approach pursued, the framework factors and this study’s low-to-high qualitative ranking approach could be fed into that process, although interdependency among the framework factors poses challenges to the use of a simple additive multiattribute model and the use of correlated input distributions would be required. A more complex multiplicative model could be considered to account for the interdependencies, but that approach adds significant complexity. For a quantitative approach, consideration would also need to be given to appropriately representing uncertainty.

In summary, this chapter describes the development of a multiattribute framework that identifies the factors that drive levels of concern for synthetic biology capabilities (the relevant outcome). The guiding objective of this approach was to identify the features of a synthetic biology capability that would affect the level of concern about a given capability being used for harm. The resulting framework is thus intended to describe the reasoning behind what is of relatively higher concern and what is of lower concern, among the capabilities considered, and why. The framework is also intended to serve as a tool that others can use to assess relative concern, albeit not in a formulaic or checklist manner, for newly emerging capabilities and to update the level of concern for existing technologies or capabilities in response to scientific and technical advances. The use of the framework to analyze specific synthetic biology capabilities is described in Chapters 4–6. Chapter 9 discusses the overall landscape of concern and presents results of the holistic assessment across the set of synthetic biology capabilities evaluated (see Figure 9-1).

4

Assessment of Concerns Related to Pathogens

The use of disease as a weapon is thought to date back to at least the Middle Ages, when the Tartars used catapults to hurl plague victims over protective walls in the city of Caffa (Wheelis, 2002). Settlers to North America presented Native Americans with blankets that had covered smallpox victims, potentially exposing this naïve population to the scourge of smallpox (Duffy, 1951). With the advent of microbiological techniques, it became possible to use specific pathogens as weapons. This capability enabled several nations, but most extensively the Soviet Union and the United States, to develop offensive biological weapons programs, which continued until they were legally prohibited by the Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on Their Destruction (known as the Biological Weapons Convention, or BWC), signed in 1972 (BWC, 1972). After the BWC was signed, the development of pathogens as weapons became the province of clandestine nation-state programs and non-state actor terrorism. One of the most high-profile uses of pathogens as weapons was the “Amerithrax” bioterror attack in 2001, in which *Bacillus anthracis* spores were sent through the U.S. Postal Service, resulting in five deaths, prophylaxis of 30,000 individuals due to potential exposures, and hundreds of millions of dollars in decontamination expenses (DOJ, 2010).

In these historical examples, naturally occurring pathogens were developed as biological weapons. Specific pathogens were selected for bioweapons development based on their ability to cause morbidity and mortality and on their ability to be converted into large-scale weapons. The age of synthetic biology raises the possibility that pathogenic bioweapons could be designed, developed, and deployed in new ways that depart from the disease-causing characteristics of a naturally occurring pathogen. First, although security protocols such as the Federal Select Agent Program (CDC/APHIS, 2017) and The Australia Group (2007), primarily in North America and Western Europe, have attempted to limit access to dangerous pathogens for many years, synthetic biology makes it possible to synthesize genomes and use those to generate, or “boot,” copies of naturally occurring organisms in the laboratory, opening new opportunities for the acquisition of existing, regulated pathogens. Second, synthetic biology techniques could be used to modify existing organisms that are not subject to limited-access regulations, potentially leading to the acquisition of desired attributes. For example, such manipulations could potentially result in pathogens that have, in comparison to the original pathogen, increased virulence; antibiotic resistance; ability to produce toxins, chemicals, or biochemicals; or ability to evade known prophylactic or therapeutic modalities. Third, synthetic biology tools could be used to synthesize and boot entirely new organisms, potentially incorporating genetic material from multiple existing organisms (Zhang et al., 2016).

This chapter analyzes these potential applications of synthetic biology related to the creation of pathogen-

based bioweapons. To assess the level of concern warranted for each capability presented in this chapter (as well as those presented in Chapters 5 and 6), the factors outlined in the report's framework for assessing vulnerabilities were considered: Usability of the Technology, Usability as a Weapon, Requirements of Actors, and Potential for Mitigation. Conclusions regarding the relative level of concern for each capability as it relates to each factor are presented in the form of a five-point scale from Low Concern to High Concern. Although all of the factors and elements identified in the framework were considered during the assessment, the discussion presented in these chapters focuses primarily on those elements deemed most salient to, or in some cases unique to, each capability. For each factor, the level of concern warranted for each capability relative to the other capabilities considered is presented at the end of the chapter along with a summary of the elements driving that relative level of concern. Conclusions regarding the relative ranking of all synthetic biology capabilities considered in the report are presented in Chapter 9.

RE-CREATING KNOWN PATHOGENS

The construction of an organism from scratch requires at least two steps: synthesis of the organism's genome and conversion of that nucleic acid into a viable organism ("booting"). Figure 4-1 illustrates these conceptual steps.

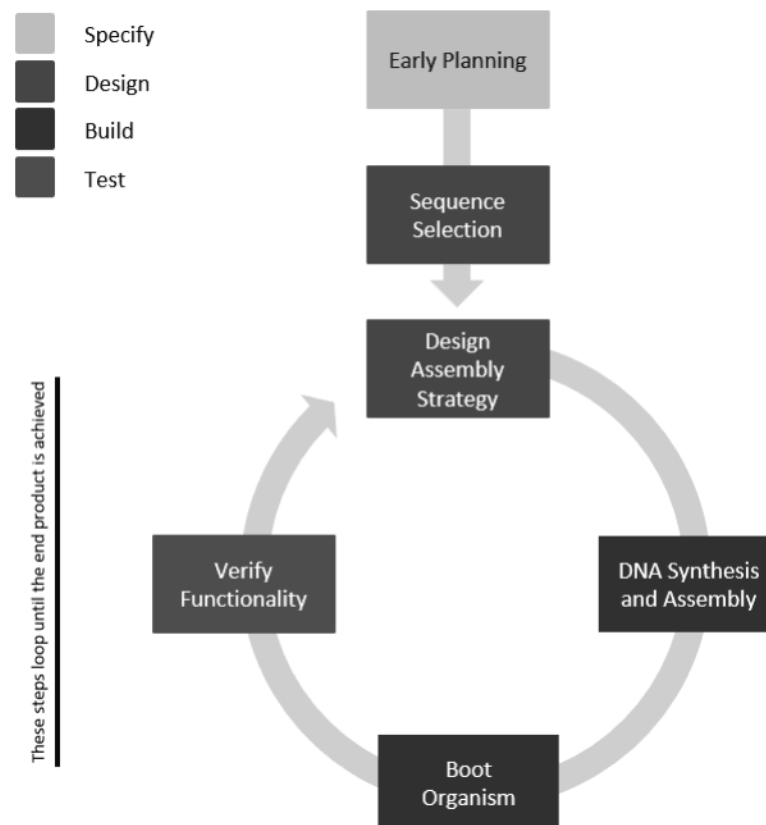


FIGURE 4-1 Activities involved in the construction of an organism from scratch. Considerations in the Design stage may include whether an exact copy of a pathogen sequence is desired, if synonymous mutations are introduced, or if a library (quasispecies) of sequences will be designed. Obtaining physical material in the Build stage may occur in the same physical location as the Design stage or may be outsourced to a commercial DNA synthesis provider. The size of the target sequence may make assembly necessary. Function of the synthesized pathogen, which may include the ability to infect and/or replicate, is determined in the Test stage.

This study assessed the potential for actors to use synthetic biology technologies to construct known, naturally occurring pathogenic organisms from scratch. Viruses and bacteria are assessed separately because of their distinct biological features. At present, construction of eukaryotic pathogens with larger genomes—such as fungi, yeast, and parasites—is considered significantly more difficult, and successes have not yet been reported.

Re-creating Known Pathogenic Viruses

Using today’s technology, the genome of almost any mammalian virus can be synthesized, and the sequences of known human viruses are readily available through public databases such as GenBank®, an annotated collection of all publicly available whole and partial DNA sequences (NCBI, 2017). The 2002 synthesis of poliovirus by Eckard Wimmer and colleagues was among the first reported syntheses of a viral genome (Wimmer, 2006). The team assembled a complementary DNA (cDNA) of the poliovirus genome (approximately 7,500 nucleotides), under the control of the phage T7 promoter, from a series of oligonucleotides with an average size of 69 bases. This cDNA was used to produce viral RNA, which was then used to program an in vitro extract to produce infectious poliovirus virions (Cello et al., 2002). Since then, larger and larger viral genomes have been generated, taking advantage of advances in the ability to synthesize longer and longer segments of DNA. Modern assembly methods have greatly expanded the scale at which DNA can be constructed, to the point that building the genome of virtually any virus—either in the form of the genome itself for a DNA virus or as a cDNA of an RNA virus that can be transcribed into the viral genome—is now possible (Wimmer et al., 2009). A notable example is the recent report of the construction of the horsepox genome (consisting of more than 200,000 base pairs) as part of an effort to develop a new smallpox vaccine (Kupferschmidt, 2017; Noyce et al., 2018). (It should be noted that while the booting of some viruses, e.g., polio, has been performed using cell-free extracts, most viruses must be booted inside cells, and some viruses, including horsepox, require the use of a helper virus in cells.)

The assessment of concerns related to re-creating known pathogenic viruses is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for re-creating known pathogenic viruses | High | Medium-high | Medium | Medium-low |

Usability of the Technology (High Concern)

Overall, the cost of producing a viral sequence and booting it is fairly low; synthesis is inexpensive and becoming more so as time passes, and cell culture facilities are not expensive to build, maintain, and operate. Therefore, since the usability of the technology is hindered only by weak barriers, the level of concern with regard to this factor is relatively high.

The Design phase of the Design-Build-Test cycle could be skipped for the synthesis of a known virus, assuming that the sequence of the genome to encode the pathogen is known. The first step of the Build phase would be to synthesize the DNA encoding the virus genome, which can either be ordered from commercial vendors or, if the actor has appropriate resources, synthesized in-house. The former approach may present a barrier because most nucleic acid synthesis companies screen for sequences of concern, such as sequences derived from pathogens on the Federal Select Agent Program Select Agents and Toxins list (CDC/APHIS, 2017). However, this barrier is weak for several reasons, including that actors need not limit themselves to viruses on the Select Agents list, industry compliance with the screening guidelines is voluntary, and oligonucleotide orders are not screened. Actors could exploit these factors or use other approaches to bypass screening, at least for viruses with smaller genomes.

Having a genome in hand is only the first step in booting a viable organism. The ease with which a virus can

be generated from its genome is largely a function of two variables: the size of the genome and the nature of the genomic nucleic acid (i.e., DNA, positive-strand RNA, or negative-strand RNA). In general, the genome must be introduced into cells in culture in which the viral genome can be replicated and assembled into infectious viral progeny. If there is no cell line in which the virus can be grown, the options become more limited. Poliovirus has been assembled completely in vitro from purified components or crude extracts (Cello et al., 2002). Although this method may become applicable to other viruses as the study of virus assembly leads to better in vitro assembly systems, such systems are currently not scalable for the production of larger quantities of virus, and eventually the actor would need to move into cell culture approaches.

Positive-strand RNA viruses, whose genomes can be directly translated by the cell to produce viral proteins, are generally easier to synthesize and boot than negative-strand RNA viruses. For positive-strand RNA viruses, the complementary DNA (cDNA) must be engineered to express an exact copy of the viral genome, including appropriate sequences at the 5' and 3' ends that govern transcription and translation, but that process is fairly straightforward. This cDNA can be transcribed in vitro to produce a viral RNA that, when transfected into cells, serves as a messenger RNA (mRNA) for production of viral replication proteins that initiate the complete viral life cycle (Kaplan et al., 1985). RNA viruses with a negative-strand genome present a slightly higher challenge to synthesize because, by definition, negative strands are not translated. For these viruses, the genome is usually introduced in the cell along with an expression vector that encodes the viral replication protein(s). Then, once the cellular RNA polymerase produces the viral RNA genome from the cDNA, the viral replication machinery can take over (Neumann et al., 1999).

Assuming that an actor can identify a cell line in which the virus can be grown, smaller viral genomes would be, in general, easier to boot, whereas large viral genomes would present a greater challenge (see Figure 4-2). Large DNA molecules must be manipulated with care to avoid fragmentation, and therefore large genomes (greater than about 30,000–50,000 base pairs) are subject to integrity constraints. However, overlapping DNA fragments are recombined readily once inside the cell, and in fact this ability to use the cell to stitch together fragments (Chinnadurai et al., 1979) was used extensively in the early days of gene therapy to produce adenovirus vectors expressing various transgenes. As the DNA of most DNA viruses is infectious, once that DNA enters

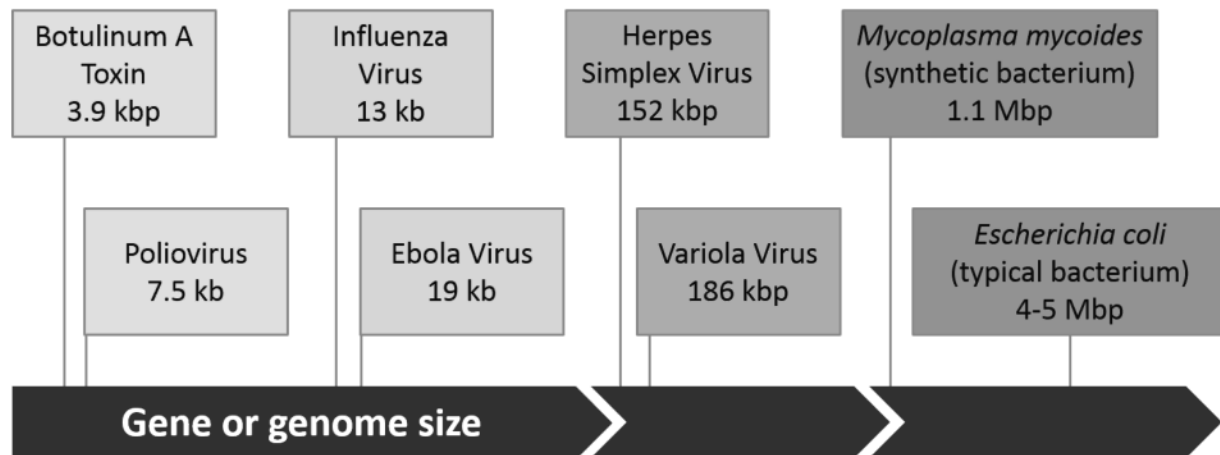


FIGURE 4-2 Relative scales of genetic information encoding familiar bacteria, viruses, and toxins. A single large toxin gene (smallest size represented in the figure, kilobase pairs) is shown in the leftmost box (lightest blue). Progressively larger genome sizes are shown in progressively darker hues moving to the right: single-stranded RNA virus genomes (kilobases), double-stranded DNA virus genomes (kilobase pairs), and bacteria (megabase pairs). The difficulty of DNA assembly and booting is partly a function of genome size and structure.

SOURCE: Adapted from John Glass, J. Craig Venter Institute.

the nucleus, the cell takes over the process of transcription and translation, ultimately leading to assembly of progeny. Poxviruses are a notable exception in that they replicate in the cytoplasm and require co-infection with a helper virus to initiate the first round of replication. The recent successful construction of the horsepox genome, which contains more than 200,000 base pairs, underscores the increasing feasibility of booting larger genomes (Kupferschmidt, 2017; Noyce et al., 2018).

Usability as a Weapon (Medium-High Concern)

Viruses have evolved to infect people and other organisms. The impact of a synthesized existing virus would be highly predictable based on knowledge of its natural behavior. The level of concern with regard to usability as a weapon spans a wide range depending on a particular virus's natural tropism, virulence, environmental stability, and other such parameters. Production scale and delivery have long been considered key barriers to using existing viruses as weapons, based on knowledge of historical offensive biological weapons programs (Guillemin, 2006; Vogel, 2012). Even today, scaling up production and delivery enough to use a synthesized existing virus as a larger-scale weapon would present substantial barriers compared to a smaller-scale attack. However, the concern level is medium-high because an actor could synthesize just a small amount of virus known to be particularly dangerous, deliver it to a small number of victims, and wait for the virus to spread as it does naturally. There are natural viruses with reproduction rates, routes of transmission, and virulence that are concerning because of the potential rapidity of spread through a targeted population after initial release or infection.

Requirements of Actors (Medium Concern)

The concern based on the requirements of actors is medium. The production of most DNA viruses would be achievable by an individual with relatively common cell culture and virus purification skills and access to basic laboratory equipment, making this scenario feasible with a relatively small organizational footprint (including, e.g., a biosafety cabinet, a cell culture incubator, centrifuge, and commonly available small equipment). Depending upon the nature of the viral genome, obtaining an RNA virus from a cDNA construct could be more or less difficult than obtaining a DNA virus. Overall, however, the level of skill and amount of resources required to produce an RNA virus is not much higher than that for a DNA virus. There are ongoing efforts to improve the nature of the cDNA clones used to produce RNA viruses (e.g., Aubry et al., 2014; Schwarz et al., 2016), but these advances tend to be incremental in nature. The J. Craig Venter Institute (JCVI) was able to develop a viable seed stock within just 3 days of learning the sequence of a new strain of influenza A virus (a negative-strand virus). Although JCVI has extensive resources and expertise that would not be available to every actor, the demonstration nonetheless underscores current capabilities regarding booting both DNA and RNA viruses.

On the other hand, one key challenge when producing some RNA viruses is the concept of quasispecies. Because viral RNA polymerases are highly error-prone, each time an RNA viral genome is copied within the cell, it generally contains one or more mutations (Lauring et al., 2012). Thus, the progeny viruses that egress from an infected cell are not a clonal population, but rather a mixture of highly related, nonidentical viruses referred to as a quasispecies. The potential genetic composition of the population, therefore, is a function of the starting sequence because any given codon can only mutate to certain other codons. Because most sequences deposited into databases are derived from recombinant clones, each of which represents a single member of the quasispecies, it is possible that the starting sequence may not generate a "wild type," fully virulent population after booting. Thus, depending on the resources and expertise available to the actor, there may be difficulties in building and testing a fully virulent RNA virus.

Potential for Mitigation (Medium-Low Concern)

The consequence management measures for attacks using re-created known pathogenic viruses would be identical to those available for the natural pathogens, including vaccines and antivirals for some agents, along with public health measures such as social distancing and isolation of sick individuals. With current approaches, it may

prove challenging to recognize and attribute such an attack because infections arising from a natural pathogen may be indistinguishable from those arising from the synthesized version. However, the same public health measures will be implemented regardless of whether the virus is synthesized or natural. While public health measures deployed to counteract natural viral outbreaks are not perfect, ongoing surveillance and containment efforts in the United States are impactful and have been effective in containing some outbreaks in recent years.

Screening commercially produced synthesized DNA sequences may be one of the only practical options to deter an attack using a re-created known pathogenic virus. The effectiveness of this approach, however, is undermined by the inherent limitations of list-based screening, the expectation that there are international companies that do not screen orders and are outside of U.S. regulatory control, the fact that oligonucleotides are not screened, and the fact that it is possible to synthesize genetic material in-house with purchased equipment.

Despite current inability to attribute and effectively prevent attacks using synthesized viruses, overall concern with regard to the potential for mitigation is medium-low owing to the existing public health measures that could be employed against an attack. However, the concern level is higher for viruses that spread rapidly and efficiently and have a short serial interval (the time between when a person is infected with a pathogen and when he or she can spread it to others).

Re-creating Known Pathogenic Bacteria

The genomes of many existing bacteria have been characterized, and the same types of DNA synthesis and booting approaches used for large viral genomes can, in theory, be applied to re-create known pathogenic bacteria. Indeed, JCVI reported the synthesis and booting of *Mycoplasma mycoides* in 2010 (Gibson et al., 2010). Other microbial genome synthesis projects are well under way, such as for *Escherichia coli* (4 million base pairs; Ostrov et al., 2016) and yeast (11 million base pairs; Mercy et al., 2017; Mitchell et al., 2017; Richardson et al., 2017; Shen et al., 2017; Wu et al., 2017; Xie et al., 2017; W. Zhang et al., 2017).

The assessment of concerns related to re-creating known pathogenic bacteria is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|--|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for re-creating known pathogenic bacteria | Low | Medium | Low | Medium-low |

Usability of the Technology (Low Concern)

It is not yet possible to successfully re-create known bacteria; therefore, the level of concern is relatively low with regard to the usability of the technology. As is the case with viruses, GenBank® is a rich source of sequence information from which to build a known bacterium. However, given that bacterial genomes are typically one to two orders of magnitude larger than most viral genomes (see Figure 4-2), bacteria present a much greater technical challenge to synthesize and boot. In the case of the JCVI synthesis (Gibson et al., 2010), a single base-pair mistake initially prevented booting of the bacteria and cost the project team months of time (JCVI, 2010). Therefore, while the Design step is straightforward, the Build component of the Design-Build-Test cycle, in particular the construction of the full genome, currently is a significant barrier. In part, this difficulty stems from the challenge of maintaining the structural integrity of the DNA itself: DNA fragments larger than 30,000 base pairs are easily fragmented when subjected to any kind of shearing, including standard laboratory pipetting, which makes them unusable for bacterial construction. To overcome this barrier in the only synthesis of known bacteria in the literature to date, the JCVI group built the bacterial genome as a yeast artificial chromosome.

Assuming the bacterial genome can be synthesized and assembled, the next step—booting—is another particularly difficult challenge, because one cannot simply add the genome to an in vitro extract and obtain a living

bacterium at the end of the reaction. Rather, the genome must be introduced into a cellular structure. The JCVI group accomplished this by transplanting their synthetic genome, propagated as a yeast artificial chromosome, into a related species of mycoplasma (Gibson et al., 2010). This transplantation approach has its own hurdles, both known (such as bacterial restriction or modification systems) and unknown. The process by which a synthetic bacterial genome may take over all necessary functions from a natural one is incompletely understood. Therefore, while obtaining the starting DNA components of a bacterial genome may be relatively straightforward from a technical point of view—they can be synthesized in-house or purchased (assuming they pass or evade Select Agents screening protocols)—the subsequent assembly steps present a substantially greater challenge than with viruses. As John Glass, leader of JCVI's Synthetic Biology and Bioenergy Group noted in a public data-gathering session during the study process, making a bacterium is "very hard and expensive."

Given that the greatest bottleneck in re-creating known pathogenic bacteria is the step that moves from DNA to functioning organism, it will be important to watch for technological advances that may facilitate genome assembly and booting. For example, the development of a method to manipulate large DNA fragments without physically damaging them could reduce the difficulty of assembly. Or if a technique were developed that allowed direct transfer of the bacterial chromosome from the yeast in which it was built into a bacterial host, this would overcome the hurdles of shearing and transplantation. However, yeasts are not known to even transfer chromosomes among themselves, except during mating; therefore, such a yeast-bacterial system would likely need to be developed from scratch if this approach was going to be pursued.

Usability as a Weapon (Medium Concern)

If a pathogenic bacterium were successfully synthesized, its properties as an infectious agent would be predictable based on the known properties of the naturally occurring bacterium. As with synthesized viruses, the level of concern therefore depends on the bacterium's natural tropism, virulence, environmental stability, and other such parameters. As with viruses, scaling up production and delivery enough to use synthesized bacteria as a weapon of mass destruction would present substantial barriers compared to a smaller-scale attack, raising many classical weaponization issues such as environmental stability during mass dispersal. Overall, the level of concern related to usability as a weapon is medium, but there is a wide range of concern with regard to different bacterial pathogens, reflecting differences in the potential for weaponization of various types of bacteria in general. For example, a bacterium that forms spores should be easier to disperse throughout, and would be more stable in the environment compared to a bacterium that does not form spores.

Requirements of Actors (Low Concern)

Making an existing bacterium from scratch currently is very difficult and requires substantial expertise and resources—significantly more resources than would be required to synthesize a known virus. Therefore, concern on this factor is relatively low. An actor would need specialized, hands-on experience working with large bacterial genomes, a level of sophistication that takes years to achieve and is currently rare. In addition, this work would require a large amount of money and a fairly long time, as evidenced by the experience of groups working in this area, such as JCVI.¹ This would likely necessitate a large organizational footprint. Thus, the capability to both construct and boot such genomes is likely to remain accessible only to large, multidisciplinary teams that have access to substantial resources (funding, equipment, diverse and well-developed skill sets) for at least the next 5 years.

Potential for Mitigation (Medium-Low Concern)

Overall, concern with regard to the potential for mitigation is medium-low due to the well-established response options that are in hand for known bacteria. In terms of consequence management, there is a wide array of antibi-

¹ The 2010 creation of the synthetic *Mycoplasma mycoides* bacterial cell by JCVI reportedly took 15 years and cost \$40 million to accomplish (see JCVI, 2010; Sleator, 2010).

otic drugs that could be used to contain attacks using bacterial pathogens (indeed, a wider array than the number of antivirals available). However, antibacterial drug resistance can be expected to limit the number of drugs that would be effective in any given case, and the re-creation of a highly virulent, antibiotic-resistant bacterium capable of aerosol transmission would pose greater concern.

In terms of prevention, it would be extremely difficult, if not impossible, to distinguish a facility being used to develop bioweapons based on synthesized pathogenic bacteria from a legitimate academic or commercial facility. The Federal Select Agent Program may provide some deterrence for these activities within the United States, although screening protocols leave many loopholes that could allow for the undetected synthesis of bacterial genome fragments for Select Agents. Also, considerations related to recognizing and attributing an attack using synthesized bacteria are identical to those for synthesized viruses; it may be quite difficult to distinguish infection by a natural pathogen from that arising from the synthesized version.

MAKING EXISTING PATHOGENS MORE DANGEROUS

The age of synthetic biology has enabled the manipulation of viruses and bacteria to alter their genotypes, and therefore their phenotypes. The gene therapy field has made engineering the tropism of viruses an active area of research, and bacteria are commonly manipulated to serve as a platform for the production of useful compounds. These same experimental approaches could be used to develop new weapons. Traits of viruses and bacteria (both pathogenic and nonpathogenic) that could potentially be modified to engineer bioweapons—along with current technological capabilities and anticipated future developments relevant to pursuing such activities—were considered in assessing the level of concern warranted for the potential use of synthetic biology to make existing pathogens more dangerous.

Making Existing Viruses More Dangerous

An actor seeking to make an existing nonpathogenic virus pathogenic or an existing pathogenic virus more dangerous or better suited for a biological attack would have multiple routes to consider. There are already some examples in the literature in which the use of biotechnology has resulted in a virus with enhanced virulence, an expanded host range, or other features that make it more pathogenic. In analyzing the level of concern warranted for this type of activity, a number of viral traits that potentially could be attempted using synthetic biology or standard techniques were considered (see Box 4-1).

The assessment of concerns related to making existing viruses more dangerous is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for making existing viruses more dangerous | Medium-low | Medium-high | Medium | Medium |

Usability of the Technology (Medium-Low Concern)

Overall, the usability of the technology required for this capability involves many barriers, leading to an assessment of medium-low concern for this factor. Although scientists have a strong understanding of viruses and their biology and can conceive of many ways to manipulate them, modifying viral characteristics intentionally using rational design remains a substantial challenge. In most cases, the viral phenotype is a result of many interrelated viral functions resulting from a diverse array of genetic networks as well as host and environmental factors. Good examples of this complex situation are found in the reviews by Herfst et al. (2017) and Plowright et al. (2017), which discuss drivers of airborne transmission and zoonotic spillover, respectively. Rarely can a specific phenotype

BOX 4-1 Viral Traits

The following are selected examples of viral traits, presented to give a sense of the range and type of traits that could theoretically be targeted for modification using biotechnology.

Altered Tropism

Tropism is the capacity of a virus to infect or damage specific cells, tissues, or species. While tropism is primarily influenced by the interaction of the viral cell attachment protein(s) with the receptor(s) present on the cell (thus determining viral entry), the larger property of tropism is determined by multiple viral and host cell factors (Heise and Virgin, 2013). Altering tropism could be used to expand the host range of an existing virus or otherwise increase a virus's ability to take hold in a targeted population.

Several studies have demonstrated the ability to alter the tropism of viruses. The avian influenza H7N9 strain has been causing isolated human infections since the initial outbreak in China in 2013, but sustained human-to-human transition has not been documented. In a recent publication, de Vries and colleagues (2017) demonstrated that only three mutational changes in the sequence of the hemagglutinin gene are sufficient to switch the virus's tropism from avian to human and support binding to human tracheal epithelial cells. However, the researchers did not perform follow-up experiments to test whether these mutations were sufficient to make an actual host range shift in the ferret model. In earlier studies with avian influenza, researchers used site-directed mutagenesis to introduce mutations into the hemagglutinin gene to allow wild-type H5N1 virus to bind to human receptors (Herfst et al., 2012). This group went on to show that as few as five mutations can lead to airborne transmissibility of H5N1 between ferrets (Linster et al., 2014).

Researchers have also used synthetic biology to alter tropism in investigations of the respiratory syndromes SARS (severe acute respiratory syndrome) and MERS (Middle East respiratory syndrome). There is considerable evidence indicating that a SARS-like virus in bats was the origin of the 2003 outbreak of SARS in humans (Li et al., 2005). The bat virus, however, does not grow in cell culture. To help elucidate the steps that may have occurred to convert bat SARS-CoV into a virus infecting humans, Becker and colleagues (2008) substituted the human SARS coronavirus receptor binding domain for the equivalent domain in the bat SARS-CoV virus, making the bat-SARS virus replication competent in cell culture and mice. Similarly, to develop a small-animal model of MERS-CoV, researchers modified both the mouse, to express a chimeric receptor, and the virus (Cockrell et al., 2016).

Enhanced Viral Replication

Enhancing viral replication could help increase the impact and spread of a virus-based bioweapon. In experiments with echovirus 7, Atkinson and colleagues (2014) demonstrated that decreasing the CpG and UpA frequencies in two 1.1- to 1.3-kilobase regions of the viral genome enhanced viral replication in susceptible cells. Conversely, increasing the CpG and UpA frequencies resulted in decreased viral replication. While it is unknown whether these results would be the same in animals—enhanced replication in cell culture does not necessarily correlate with enhanced replication in vivo, and in fact, the reverse is sometimes the case—an actor with sufficient time and resources may be able to generate variants empirically and passage them in a susceptible host to select a variant with enhanced replication ability.

Enhanced Virulence

Virulence measures the relative capacity of a virus to cause actual disease in a host, rather than just infection. Virulence represents the combined effect of multiple genes and determinants

continued

BOX 4-1 Continued

that play specific roles in specific settings in vivo (Heise and Virgin, 2013). In the best-known example of an engineered virus resulting in enhanced virulence, Jackson and colleagues (2001) engineered ectromelia virus (mousepox), a member of the *Orthopoxvirus* genus and a natural pathogen of mice, to express mouse interleukin-4 (IL-4), with the goal of producing a contraceptive vaccine to control the mouse overpopulation. In the mouse model, the recombinant virus was shown to suppress primary antiviral cell-mediated immune responses and overcome preexisting immunity. It is also conceivable that actors would seek to manipulate a virus so that it causes disease by different mechanisms than a natural virus might, such as by manipulating neurobiology or altering the host microbiome.

Ability to Evade Immunity

At the root of the increased virulence demonstrated in the mousepox experiments (described under Enhanced Virulence, above) was the recombinant virus's capability to evade immunity. This points to another potential route for actors seeking to produce bioweapons: the development of viruses designed to anticipate and evade the immune response or even to overcome vaccine-based immunity. Detection of viral pathogens by the innate immune system leads to the induction of antiviral mechanisms that are mostly mediated by type-1 interferons. This primary response then leads to the activation of the adaptive immune response that is more directed, antigen-specific, and longer lasting (Iwasaki and Medzhitov, 2013). Many viruses have countermeasures to subvert the innate immune response including interferon-induced antiviral activity (see Chan and Gack, 2016, for a review). It may be possible to express one or more antagonists of these antiviral activities in a pathogen that does not already have that particular antagonist. In this way, the arsenal of activities that a virus uses to evade the innate immune response would be expanded and virulence may be enhanced.

The creation of chimeric viruses developed by genetically substituting capsid genes has been well documented (see Guenther et al., 2014, for a review). These viruses have mainly been developed in the context of, for example, improving adenovirus vectors to target specific tissues and as an approach to circumventing preexisting viral immunity that may limit the use of viral gene therapy vectors (Roberts et al., 2006). It is conceivable that the latter approach could be used to develop a chimeric viral vector expressing a toxin gene targeted to a particular tissue and used in a population with preexisting immunity to the vector virus. The molecular determinants of targeting are poorly understood, however, and these approaches generally require significant trial and error to be successful.

Ability to Evade Detection

Some modifications could result in a virus that would be difficult to detect using current outbreak response approaches. The most commonly used methods of laboratory identification of viruses are based on real-time polymerase chain reaction assays in which specific primers and fluorescently labeled probes are designed to bind to conserved and unique regions of the viral DNA or cDNA. Nontargeted methods of detection include array-based assays and next-generation sequencing, but these are not yet in wide use in clinical and commercial laboratories. Cell culture methods are rapidly disappearing from use. Mutations that target the primer binding sites could therefore result in a virus that is not recognizable.

Ability to Resist Therapeutics

Actors could seek to develop viruses capable of resisting available therapeutics, though the necessity of this approach would depend on whether effective therapeutics exist. Despite the availability of successful antiviral agents such as those used to counter HIV (human immunodeficiency virus), herpes viruses,

influenza viruses, and HCV (hepatitis C virus), there are no specific antiviral drugs for the vast majority of viruses. Even where antivirals exist, the development of resistance to these drugs is almost inevitable unless the rate of replication of the virus in the presence of the drug can be completely inhibited or, alternatively, if multiple drugs are used in combination against different viral targets (Coen and Richman, 2013). For example, newer antivirals based on immune inhibition, such as the ZMapp therapeutic, are a mixture of three humanized monoclonal antibodies developed against Ebola virus and have shown survival benefits in nonhuman primates experimentally infected with the virus (Pettitt et al., 2013). A randomized, controlled trial in humans appeared to show beneficial effects but did not meet the prespecified statistical threshold for efficacy (Davey et al., 2016).

Enhanced Transmissibility

Airborne transmission of pathogens occurs through aerosolization and droplets. Airborne transmissibility determines the distance over which the virus may travel, and the determinants of this property are complex and dependent on multiple host and viral factors (Herfst et al., 2017). In a follow-up to the H5N1 experiments described under Altered Tropism (above), the mutated virus was sequentially passaged in ferrets to force natural selection of heterogeneous viral mixtures and, after 10 passages, naïve recipient ferrets were exposed to the infected ferrets in an adjacent cage without direct contact. Three of four recipient ferrets became infected, demonstrating that selection had occurred for airborne transmissibility of the virus (Herfst et al., 2017). In another study, Imai and colleagues (2012) constructed a reassortant virus possessing the hemagglutinin from an H5N1 virus and seven gene segments from a 2009 H1N1 virus. After passaging through ferrets, a mutant of this reassortant was obtained that had four mutations in the hemagglutinin gene and was capable of respiratory droplet transmission in ferrets. This work demonstrated that a mammalian transmission phenotype could be conferred to highly pathogenic H5N1 influenza.

Enhanced Stability

The stability of a virus outside the host is influenced by multiple environmental factors including temperature, ultraviolet radiation, relative humidity, and air movement, as well as the structure of the pathogen itself. Enveloped viruses are generally less stable outside the host than non-enveloped viruses (Polozov et al., 2008; Herfst et al., 2017). Although it would be impossible to convert an enveloped virus to a non-enveloped virus because addition of the envelope is tightly coupled to specific features of the replication cycle, it may be possible to alter other features of a virus to enhance its stability for weaponization and mass dispersal.

Reactivation of "Dormant" Virus

It may be possible to use chemical or biological means to reactivate latent or persistent viruses. Such an attack could be targeted based on whatever endogenous mix of pathogens already exists in an individual or population. For example, some viruses, like HCV, cause chronic infections whose clinical symptoms do not appear until late in life; developing a chemical or biological trigger to accelerate the pathogenesis of such a virus is a possibility. It may even be possible to recombine a modern virus that has little pathogenicity and spreads widely with an earlier, perhaps more deadly, endogenous variant.

Lower immunity in hematopoietic stem cell transplant patients has been shown to result in widespread viral reactivation, sometimes life-threatening (Cavallo et al., 2013), underscoring the potential impact of such approaches. Research focused on coaxing HIV out of latent reservoirs in order to completely cure the infection, the so-called "shock and kill" strategy (Shirakawa et al., 2013), could further advance potential dual-use research in this area.

be attributed to a single gene, or an altered phenotype to a specific mutation. Furthermore, the determinants of tropism, transmissibility, and other properties are often not well understood or predictable. Many of the research advances achieved to date have involved significant trial and error (e.g., gene therapy vector tropism modifications [Nicklin and Baker, 2002]), inadvertent findings (e.g., the outcomes of IL-4 expression in ectromelia virus [Jackson et al., 2001]), or directed evolution (e.g., experiments altering transmissibility of avian influenza virus [Herfst, 2012; Imai et al., 2012]). How these alterations would affect the behavior of these viruses in the human population is difficult to assess because of limited knowledge regarding how genotype would translate to phenotype, but a successful introduction of such a modified virus into humans could have dire consequences. Although this knowledge gap of how to engineer complex viral traits is likely to limit the ability to engineer viruses for enhanced bioweapons currently, it will be important to monitor for developments that significantly increase the ability to relate genotype to phenotype—the knowledge of determinants of complex viral traits and how to engineer pathways to produce them.

An added barrier is that introducing mutations into a viral genome almost invariably results in an attenuated (i.e., less pathogenic) virus (Holmes, 2003; Lauring et al., 2012), because there are constraints on viral genome organization. The introduction of mutations has been the classical method of making many effective live attenuated vaccines, including those for measles and yellow fever, as well as the Sabin poliovirus vaccine strain (Sabin, 1985). The mutation(s) in these examples were introduced in a nondirected manner by passage in cell culture and resulted in phenotypic changes that lessened the virus's ability to cause a harmful infection. An exception to this assessment of medium-low concern, however, would be the introduction of antiviral resistance. It is more feasible to introduce mutations that allow resistance to antivirals without causing attenuation, because the exact point mutations responsible for drug resistance are often known and generally do not lead to significant attenuation.

The majority of alterations in a viral genome can be performed with standard recombinant DNA technology methods and do not require advanced synthetic biology techniques. One exception is the multiple substitutions required to change the frequency of particular bases to make synonymous mutations at multiple positions. Achieving this would be much simpler with the large pieces of DNA that synthetic biology technologies assist in producing, as well as synthetic biology tools that allow for the introduction of mutations in a directed manner and the application of many mutations simultaneously. For example, researchers are now using synthetic biology to introduce many synonymous mutations (including alterations in a DNA or RNA sequence that do not change the protein amino acid sequence), in an effort to make live attenuated viral vaccines that have better genomic stability (Wimmer et al., 2009; Martinez et al., 2016).

Given the precision required and the limitations of rational design, an alternative approach would be to use combinatorial libraries, high-throughput screening, or directed evolution to test many candidate modifications. For example, viruses could potentially be tailored to evade specific immune responses by using computational modeling, high-throughput screening, or directed evolution to escape the most likely or most capable antibodies or T-cell receptors, provided that immune-dominant epitopes on a pathogen are known. However, even this approach would be constrained to some extent by the amount of available information regarding the determinants of the target phenotype and potentially by the current size limits of combinatorial libraries. It is not possible to test an infinite number of variations, although with available technologies a well-resourced actor would be capable of testing quite a lot.

Finally, in addition to developing the variants to test, it is necessary to boot the recombinant genome in a cell line. Depending on the virus, this booting step can present a significant barrier, and booting imposes additional limits on the number of variants that can feasibly be tested.

Usability as a Weapon (Medium-High Concern)

Because viruses have certain characteristics consistent with use as a weapon, and because the modification of the virus may enhance those characteristics, the concern is medium-high for this factor. Just as the types of manipulations required to alter the phenotype of a virus are difficult to predict, how a modified virus will behave when introduced into the human host is also difficult to anticipate. In addition, the tendency for alterations to attenuate viruses may serve as a “natural” mitigating factor and reduce the effectiveness of a bioweapon produced

in this way. Testing modified viruses may also present a barrier (unless the actor is willing to test in humans). For example, animal models do not always predict how a virus will behave in humans. It has been argued that avian influenza virus transmission in ferrets does not mean with certainty that those viruses would also transmit from human to human via an airborne route (Racaniello, 2012; Lipsitch, 2014; Wain-Hobson, 2014), but as noted above, if an engineered virus does acquire this property, the dynamics of weapons use change.

If modifications are pursued with the intention of making the virus more dangerous in some way, the scope of casualty for an attack using a modified virus could be larger than an attack using a natural virus. If the modifications are intended to make the virus easier to produce or deliver, the resulting virus may bypass some of the classical barriers to weaponization, such as environmental stability during mass dispersal. Otherwise, a modified virus would present many of the same weaponization opportunities and challenges as those detailed for the re-creation of a known pathogenic virus.

Requirements of Actors (Medium Concern)

Modifying a virus would require excellent molecular biology skills and advanced knowledge of the field. Understanding and being able to verify the product therefore imposes an expertise barrier to successfully manipulating viral phenotypes. In general, however, the resources and organizational footprint required would be moderate, similar to those required for re-creating a known pathogenic virus. Therefore, there is a medium level of concern with regard to this factor.

Potential for Mitigation (Medium Concern)

Existing tools for mitigation, such as public health systems and antivirals, may be effective against a modified virus. However, in general, they would be expected to be less effective against modified viruses than against the naturally occurring ones for which they are designed, leading to a medium level of concern for this factor. In particular, available medical countermeasures may be ill-suited against viruses with modifications designed to confer antiviral resistance or to alter the ability of the virus to be recognized by the immune system. Diagnostic approaches using sequencing would be effective for identifying a modified virus as being laboratory-derived in the vast majority of cases (antiviral resistance being one notable exception), but it is unclear whether that capability would effectively facilitate attribution. Although the overall level of concern for this capability is medium with regard to the potential for mitigation, the concern level is higher for viruses with pandemic potential, such as influenza, for which a modified virus could present significant challenges in terms of measures to limit spread or reduce impact.

Making Existing Bacteria More Dangerous

As with viruses, an actor seeking to make an existing nonpathogenic bacterium pathogenic or to make an existing bacterial pathogen more dangerous would have many potential routes to consider. In analyzing the level of concern warranted for this type of activity, a number of modifications to existing pathogenic or nonpathogenic bacteria that potentially could be attempted using biotechnology were considered. Box 4-2 notes some of the ways in which such activities might differ in the context of bacteria compared to viruses.

The assessment of concerns related to making existing bacteria more dangerous is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|--|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for making existing bacteria more dangerous | High | Medium | Medium | Medium |

BOX 4-2

Bacterial Traits

The following are selected examples of bacterial traits, presented to give a sense of the range and type of traits that could theoretically be targeted for modification using biotechnology. This box focuses on how modifying traits in bacteria might differ from modifying analogous traits in viruses, described in Box 4-1.

Altered Tropism

Unlike viruses, which are exclusively intracellular pathogens, bacterial pathogens can be either intracellular or extracellular. Generally, extracellular pathogens are relatively environmentally stable and good at adapting to their environment. Even those that are not spore-forming often have the capacity to replicate and cause damage in multiple tissues and cell types and in different locations in the body. Given their environmental stability, they are difficult to eradicate and may not require host-to-host contact for transmission. Intracellular bacteria, like viruses, rely on host cell nutrients and are often able to evade the host immune system (Finlay and McFadden, 2006). Intracellular pathogens are usually transmitted via direct contact or aerosol transmission. Both intracellular and extracellular pathogens rely on adhesins and colonizing factors, which facilitate contact with host target cells, confer resistance to leukocyte attack, and are significant virulence factors (Ribet and Cossart, 2015).

Enhanced Virulence

Many factors influence bacterial virulence and could potentially be targeted for modification. The primary mechanisms of bacterial pathogenesis include host target cell death (Böhme and Rudel, 2009), whether by cell lysis (resulting either from the multiplication of intracellular pathogens or as a result of the action of bacterial toxins) or by induction of apoptosis (programmed cell death); mechanical perturbations of host physiology (e.g., blockage of circulatory or respiratory passages due to the size or number of invading bacterium or as a result of mucous production); host cell damage resulting from the host immune response to the bacterial infection; and the action of bacterial toxins. The effects of cell death depend upon the host cells involved and are influenced by the bacterial burden introduced, the route of infection, complicating symptoms induced by host immune response, and the rapidity of the infection process. Colonization potential is influenced by the ability of some pathogenic bacteria (e.g., *Shigella*) to trigger premature or unscheduled apoptosis in the host cells they infect (Gao and Kwaik, 2000); the initial phase of this process involves the introduction of enzymatically driven damage to host cell DNA followed by massive disturbances in cell integrity and cell death. Another significant virulence factor is the ability of some bacteria (e.g., *Bacillus anthracis*) to form capsules consisting of polysaccharides and amino acids (Cress et al., 2014). Capsules prevent bacteria from being phagocytized by neutrophils and macrophages. Other virulence factors include invasion factors, which are usually encoded chromosomally but may also be plasmid-borne, and siderophores, iron-binding factors that allow bacteria to compete with host cells for iron acquisition (Quenee et al., 2012).

Enhanced Toxin Production

Many bacterial pathogens cause damage to host cells and tissues through the production of toxins. These toxins take two forms: exotoxins and endotoxins. Exotoxins are relatively unstable, highly antigenic proteins that are secreted into host body fluids. Some exotoxins are bound to the bacterial cell wall following their synthesis and are released upon lysis of the invading bacterium (Sastalla et al., 2016). Often highly toxic, exotoxins are produced by both Gram-positive and Gram-negative bacteria. Some exotoxins can act only on certain cell types whereas others affect a broad spectrum of cells and tissues. Some bacterial pathogens make only a single toxin (e.g., cholera, diphtheria, tetanus, botulism) whereas others

can synthesize two or more distinct toxins (e.g., *Staphylococcus*, *Streptococcus*). Antitoxin antibodies to exotoxins are usually made rapidly by the host. The genetic determinants of exotoxins are often found on extrachromosomal elements, usually plasmids or bacteriophages.

Endotoxins, on the other hand, are relatively stable, lipopolysaccharide components of the outer membrane of some Gram-negative bacteria that can act as toxins under certain circumstances (Zivot and Hoffman, 1995). Lipid A appears to be the toxic component, which can act while in the intact bacteria expressing it. Endotoxins are generally weakly immunogenic, eliciting fever in the host. They can cause hypotension due to increased vascular permeability accompanied by vasodilation, which can in turn result in shock. The genetic determinants for endotoxins are chromosomal.

Actors could potentially seek to modify bacteria to enhance their natural toxin production or introduce toxin production into a bacterium that does not naturally produce toxins. Such approaches are further discussed in Chapter 5.

Ability to Evade Immunity

As with viruses, it is possible to engineer bacteria to anticipate or evade the immune response.

Ability to Evade Detection

As with viruses, the most commonly used methods of laboratory identification of bacteria are based on real-time polymerase chain reaction (PCR) assays in which specific primers and fluorescently labeled probes are designed to bind to conserved and unique regions of the bacterial chromosomal or extrachromosomal DNA. Another widely used method in clinical microbiology laboratories is MALDI-ToF (matrix-assisted laser desorption/ionization time-of-flight), a method of ionizing large molecules and identifying them by mass spectrometry in comparison to reference standards. Nontargeted methods of detection such as array-based assays and next-generation sequencing are available but are not yet in wide use in clinical and commercial laboratories. Culture methods are rapidly disappearing from use (Carleton and Gerner-Smidt, 2016).

Ability to Resist Therapeutics

In contrast to the relatively small number of antivirals, there are many antibacterial agents available that are capable of acting against a wide variety of bacterial pathogens. However, bacteria can be intrinsically resistant to antibiotics, or can acquire resistance via chromosomal mutation and horizontal gene transfer. There are three main mechanisms of antibiotic resistance (Blair et al., 2015). First, the bacterium can prevent the antibiotic from accessing its target, either through reduced permeability of the antibiotic through the cell wall or membrane complex or through increased efflux of the antibiotic back out of the organism and away from its target. Second, the antibiotic target can be altered through genetic mutation, causing the target to become modified or protected. Finally, antibiotic resistance can be acquired by direct modification of the antibiotic itself, either by inactivation by antibiotic hydrolysis or by way of inactivation due to a chemical modification. These mechanisms are well studied and could potentially be adapted for the purposeful creation of antibiotic-resistant pathogenic bacteria.

Enhanced Transmissibility

As with viruses, the property of airborne transmission in bacteria is complex and dependent on multiple host and pathogen factors, in particular environmental stability and tissue tropism. Extracellular bacterial pathogens are extremely adaptable to environmental challenges and may not require host-to-host contact for transmission, making these pathogens difficult to eradicate. In addition, many bacterial pathogens that replicate extracellularly are capable of causing damage to different cells and tissue types. On the other

continued

BOX 4-2 Continued

hand, many intracellular bacterial pathogens are communicable (i.e., capable of host-to-host transmission), facilitating rapid spread within a community and thus presenting a greater capacity to threaten public health.

Enhanced Stability

The environmental stability of a bacterium depends on its physiology and life cycle. Generally, because of the composition and structure of cell walls, Gram-positive bacteria are more environmentally stable than Gram-negative bacteria. In addition, when subjected to harsh environmental conditions such as desiccation, some Gram-positive bacteria form spores capable of remaining viable in the environment for decades, albeit in a metabolically dormant state. For example, spores of *Bacillus anthracis* can remain viable in the environment for up to a century (Friedmann, 1994; Repin et al., 2007; Revich and Podolnaya, 2011) and constitute the infectious form of this pathogen (with vegetative forms not being infectious). Actors may find it advantageous to engineer bacterial cell walls to more closely resemble Gram-positive organisms to enhance survival during aerosol dissemination and allow the agent to remain viable and available to infect the target host for extended periods of time.

Usability of the Technology (High Concern)

Generally speaking, the technology requirements for making existing bacteria more dangerous are relatively low, which leads to a relatively high level of concern for this factor. Although it is technically difficult to design and build bacteria from scratch, altering existing bacteria is relatively easy with molecular and genetic approaches. These capabilities make the Design phase of the Design-Build-Test cycle relatively straightforward, especially if the desired trait is conferred through a well-elucidated gene or pathway, such as known genes for antibiotic resistance or toxin production. In terms of the Build step, there are well-established techniques to insert, delete, or change existing genes (Selle and Barrangou, 2015; Wang et al., 2016; H. Zhang et al., 2017). Making such modifications does not necessarily require synthetic biology approaches, though such technologies can enhance the process. Some bacterial species are easier to manipulate genetically than others. In general, this step is easier if the genetic changes are smaller in size or fewer in number and more difficult for larger or more extensive modifications. In addition, if a desired pathogen has a close nonpathogenic relative, a researcher could splice relevant portions of the pathogen's genome into the genome of the relative.

In general, it is easier to manipulate bacteria than viruses. In part, this is due to the relative sizes of bacterial versus viral genomes; for viruses there are fitness pressures and constraints on genome packaging to keep the genome smaller, thus tending to attenuate modifications over time. Modifications are more likely to persist in a bacterial genome because those genomes are genetically more stable. In viruses, enhancement of one phenotype often results in diminution of another, a factor that would likely be difficult to overcome in viruses but presents less of a barrier when modifying bacteria.

Some types of bacterial modifications would be easier to achieve than others; engineering bacterial traits that are complex requires greater knowledge of trait determinants and how to engineer pathways to produce them. On the more difficult end of the spectrum is altering tropism, which involves the complex interplay of a multitude of bacterial genes that are fundamental to the physiology of a specific bacterium (Pan et al., 2014). Tropism in bacteria is less likely to be alterable using synthetic biology approaches compared to tropism in viruses; however, there are routes that could be pursued. Both intracellular and extracellular bacterial pathogens rely on adhesins and colonizing factors to facilitate contact with host target cells (Ribet and Cossart, 2015). It may be feasible to use synthetic biology technologies and big data analytical capabilities to engineer and express novel adherin or colonizing factor analogues of these bacterial proteins and introduce them either by encoding them on episomes or integrating them into the chromosome. Given the complexity of the host-pathogen interaction, transmissibility and communicability of bacterial pathogens in humans would also be difficult to confer or alter. In a similar vein,

it would be challenging to manipulate a bacterial pathogen to acquire efficient airborne transmission. Among other characteristics, the pathogen's success would depend on environmental stability, which is intrinsic to its physiology and life cycle. It is not yet technically possible to alter a bacterial pathogen's environmental stability in a fundamental way, such as by converting a Gram-negative bacterium to Gram-positive or a non-spore-forming bacterium to a spore-forming bacterium. That said, synthetic biology approaches would have greater likelihood of success in this realm than would standard molecular biology approaches.

On the other hand, bacterial toxins, both endotoxins and exotoxins, are clearly significant virulence factors that can likely be readily modified or designed based upon data analysis. Given that endotoxins are chromosomally expressed and are intrinsic to the physiology of the bacterium in question, an actor would likely need to use a combination of synthetic biology and standard molecular biology approaches to modify existing endotoxins or create new ones. In addition, it is relatively trivial to confer resistance to antimicrobial drugs via standard molecular biology technologies (as demonstrated by the fact that it was done many years ago [Steinmetz and Richter, 1994]), and synthetic biology approaches would further enable targeted mutations to create a drug resistance phenotype.

Usability as a Weapon (Medium Concern)

The weaponization potential for making a bacterial pathogen more dangerous is, overall, of medium concern. Historically, scale-up and environmental stability have been key barriers to the weaponization of bacteria. Synthetic biology does not drastically change this equation. Despite a sophisticated understanding of some traits, such as antibiotic resistance and toxin production, knowledge is still limited for traits relevant to production and delivery of bacteria as a bioweapon, as noted under Usability of the Technology, above.

Requirements of Actors (Medium Concern)

The expertise required to design genetic modifications to affect bacterial traits varies widely depending on the nature of the modification (e.g., those that change the bacterium's biology in a new way would be more challenging) and the amount of available information about the genes involved (e.g., those involved in toxin production and antibiotic resistance are fairly well elucidated and would thus be accessible to someone with less expertise). Thus, as more information is published relevant to more traits, the level of expertise required to design modifications to those traits is reduced. Based on the current state of knowledge, this factor poses a medium level of concern.

Making the actual modifications would require classical molecular biology expertise and experience in bacterial genetic approaches, but does not necessarily require training in advanced synthetic biology techniques.

Potential for Mitigation (Medium Concern)

The current concern level for this factor is medium. As discussed in the context of re-creating known pathogens, the Select Agents list and voluntary screening guidelines are not likely to be sufficient to deter or prevent the development of modified bacterial pathogens. In terms of consequence management, one fundamental difference between responding to a naturally occurring new organism that has unique characteristics and responding to a modified bacterial pathogen that is a purposefully deployed biological weapon is a calculating adversary. Although public health system components such as the National Syndromic Surveillance Program (NSSP) of the U.S. Centers for Disease Control and Prevention may indeed be well suited to detecting and containing new naturally occurring bacterial threats, an engineered organism resistant to antibiotics will challenge the ability of public health systems to contain and respond to such a pathogen. Thus, consequence management capabilities would be less effective in the face of bacterial pathogens engineered specifically to evade them, such as through resistance to vaccines or antibiotics.

CREATING NEW PATHOGENS

A major aspiration within the field of synthetic biology is the design and creation of new organisms with beneficial uses. In the context of bioweapons, the possibility that this aspiration may potentially be directed toward producing pathogens that are entirely new was considered. In contrast with the discussion of modifying existing pathogens, the term “new” is used here to describe novel combinations of genetic parts from multiple organisms for which the product is not recognizable as primarily from one source. This can include genetic parts designed computationally with no near relative in the natural world. The resulting range of potential bioweapons in this category is extremely broad but serves to illustrate the more challenging applications that may be possible at some point in the future.

One example of a new pathogen would be a virus constructed from parts of many different natural viruses. This mix-and-match approach might be used to combine the replication properties of one virus, the stability of another, and the host-tissue tropism of a third, for example. A variety of experimental approaches would be applicable to this goal. Directed-evolution approaches could be used to sample random combinations of viral DNA parts; while each individual combination would have a small chance of success, sampling a very large number of combinations would increase the chances of success. More explicit design approaches might be to develop software to model and predict the properties of specific designs, which would then be built, tested, and improved through multiple iterations of the Design-Build-Test cycle. As discussed under Making Existing Viruses More Dangerous, however, even simple changes to existing viruses can produce drastic deficiencies in key viral properties, making any such effort especially difficult. Nonetheless, work involving recomposing the structure of a bacteriophage genome into modular pieces (Chan et al., 2005) suggests that radical new combinations of viral sequences may be viable, although tools to design viruses with high confidence of success are currently lacking.

A different example of a new pathogen would be one based on synthetic “genetic circuits” (described in Appendix A). A major pursuit within synthetic biology is the capability to arbitrarily program specific functions using genetic material. These efforts are exemplified by the engineering of DNA-encoded programs, relying heavily on concepts derived from information theory and computer science, such as constructing logic gates from individual switching functions. Importantly, the genetic material encoding those functions can in principle come from anywhere—from any branch of the tree of life or from an entirely new DNA sequence that has never been observed in nature. The designs for genetic circuits have greatly increased in complexity over time (see Toman et al., 1985, for an early example) through increased reliance on component abstractions and standardization. Figure 4-3 shows a recent example of software developed to enable such advanced designs in general, but not specifically in the context of pathogens.

Although a number of genetic circuits have been designed to function in human cell lines in culture, applications using genetic circuits in the human body are still in their infancy (Lim and June, 2017). The potential for using such technology to cause harm in the human body is thus a subject of broad speculation. Novel circuits could (in theory) be used to convert a healthy cell into a cancerous one or to provoke an autoimmune response. Such circuits might be designed to act on the host DNA using engineered factors that turn host genes on or off, such as at the level of transcription or translation. A variety of mechanisms have been demonstrated for such general-purpose switching, including the use of natural or artificial microRNA molecules and the use of CRISPR/dCas9-type programmable gene repression or activation (Luo et al., 2015). Importantly, these are examples of mechanisms that have displayed a high degree of programmability in terms of which host DNA sequences can be targeted. In a similar vein, the potential programmability of genetic effectors may also lead to genetic circuits that sense and compute based on the state or type of cell (Weiss et al., 2003) or even specific genetic identity. In some cases, genetic circuits could be delivered to a small number of host cells using nonreplicating delivery mechanisms, which could be either virus-derived, such as those used in some gene therapies (see Chapter 7, Gene Therapy), or based on nonbiological materials.

At the extreme end of difficulty (and feasibility) lies the engineering of life forms that are particularly dissimilar from known life on this planet. “Xenobiology” (described in Appendix A) offers some possibilities—for example, a bacterium employing a different combination of deoxyribonucleotides and ribonucleotides to encode

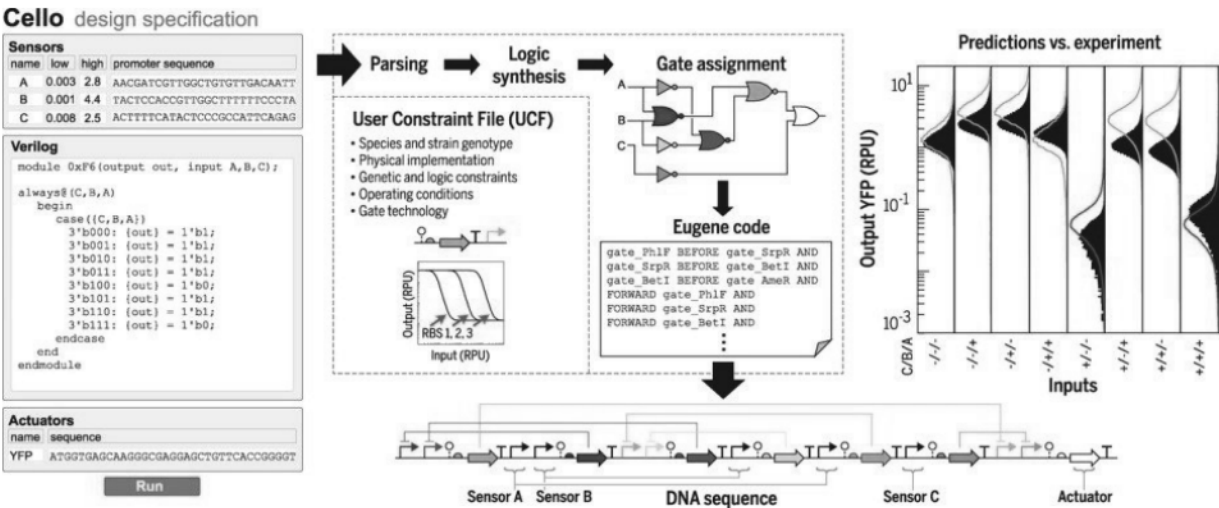


FIGURE 4-3 Illustration of genetic circuit engineering facilitated by a software environment that couples circuit specification and design to predictive models of circuit function. NOTE: Genetic circuits are a common staple for work in synthetic biology and allow users to combine multiple functions from the broad categories of sense, compute, and actuate. SOURCE: Nielsen et al., 2016.

its genetic information (Y. Zhang et al., 2017). There is a wide range of expert opinion as to the long-term plausibility of such efforts.

The assessment of concerns related to creating new pathogens is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for creating new pathogens | Low | Medium-high | Low | Medium-high |

Usability of the Technology (Low Concern)

Because the creation of new pathogens faces multiple major knowledge and technical barriers, including knowledge regarding minimal requirements for virus and bacteria viability and the constraints on viral organization discussed above, the level of concern for this factor is very low at present. However, this is a clear example of an area that warrants ongoing attention. If the technical barriers can be overcome in the future, the level of concern would increase substantially. For example, the recent engineering of a designed nucleocapsid (a protein structure capable of packaging its own genetic material, reminiscent of a virus [Butterfield et al., 2017]) demonstrates how mimicking some pathogen-like functions may be achieved without relying on pathogen-derived DNA. Nevertheless, such work falls far short of the extensive engineering required for producing a truly new viral pathogen. While

packaging genetic material is one essential viral function, additional barriers exist in engineering efficient host or tissue targeting, cellular entry, genome replication, and viral particle maturation, budding, or release. Optimizing all of these functions to work effectively in concert presents an additional difficulty. Reliably engineering a brand new virus to cause specific symptoms in the host is likely to be even more challenging.

Usability as a Weapon (Medium-High Concern)

The level of concern related to usability as a weapon is medium-high, primarily due to two factors. First, it may be possible to create pathogens with features not seen before. Such features could include, for example, the ability to target specific tissues or cell types using genetic logic, or the ability to produce aberrant neurological effects. Similarly, such pathogens could employ novel timing mechanisms, creating a delay between the time of exposure and the onset of symptoms. Second, in theory, pathogens designed from scratch may have a greater ability to cause harm because humans may not have been exposed to similar pathogens previously, and therefore may be immunologically naïve.

Requirements of Actors (Low Concern)

Design, construction, and testing of a completely novel pathogen requires capabilities that have not yet been demonstrated. While this capability is extremely broad in terms of the specific types and features of a pathogen that could be created, the high degree of expected technical difficulty leads to an overall low level of concern in terms of the requirements of actors. Furthermore, the high uncertainty that such ambitious projects would yield the desired result in itself may lead actors away from such a path toward more reliably fruitful efforts. In general, one would expect that such ambitious, envelope-pushing projects would require well-resourced teams with deep expertise in several different technologies. A successful project would also be expected to require advanced design skills and tools, in particular software platforms that enable modeling and prediction of a pathogen's properties, including host-pathogen interactions. Furthermore, navigating this uncharted territory would in general require many iterations of the Design-Build-Test cycle, with extensive testing needed during development. Thus, successfully designing and deploying a new pathogen would likely require a team of actors with significant time, money, and other resources to invest in the process and a permanent, well-equipped facility (as opposed to a mobile or makeshift laboratory).

Potential for Mitigation (Medium-High Concern)

A completely novel engineered pathogen would have the potential to frustrate existing mitigation approaches in multiple ways, leading to a medium-high level of concern for this factor. First, attempts to identify the pathogen through molecular methods—such as PCR, sequencing, or the enzyme-linked immunosorbent assay (ELISA)—would be hampered because the pathogen would not produce results that match cleanly to known pathogens. (Indeed, in some cases one could imagine partial matches to multiple pathogens.) However, analysis of the genetic sequence of the new pathogen would likely indicate that a novel biological entity is present, providing important information. Second, symptoms of the new pathogen could mislead initial attempts at diagnosis, where common pathogens would be suspected first. Third, even if the agent is identified, correct treatment choices for the new pathogen would be uncertain. However, treatment measures taken that are common across a variety of ailments (i.e., anti-inflammatory drugs, rest, fluids) might still be germane and of some effectiveness because such approaches are tied not just to the specific features of a given pathogen, but to general classes of symptoms in human disease (e.g., fevers, swelling, congestion, inflammation).

SUMMARY

- Known pathogens can be re-created. The difficulty of this re-creation increases with the size of the genome.
- Engineering viruses to make them more pathogenic is possible. Design would be challenging because of knowledge limitations and because changes are generally detrimental to viruses; however, these challenges could potentially be addressed by building and testing many variations until a more pathogenic virus emerges.
- Bacteria can be engineered with current technology, and the engineering of bacteria with characteristics such as multidrug resistance is an area of near-term concern.
- With regard to making new pathogens, the difficulty increases as the distance from natural pathogens increases.

Humans have used pathogens as tools of war for centuries. Modern biotechnology has opened new opportunities for creating bioweapons, and synthetic biology further enhances and expands these opportunities. This report examined current capabilities and expected future developments related to re-creating known pathogenic viruses and bacteria, modifying existing nonpathogenic and pathogenic viruses and bacteria, and the potential creation of entirely new pathogenic agents.

The possibility of re-creating known pathogenic viruses poses a relatively high level of concern. This concern is driven largely by the technical ease of synthesizing viruses (especially those with smaller genomes) and known pathogenicity of existing viruses (thus making them potentially reliable bioweapons). However, because current mitigation approaches were designed to counter natural viruses, they would be reasonably well equipped to mitigate synthetic versions of known viruses. Looking forward, it will be important to monitor technological advancements that make it easier to synthesize larger and larger viruses, which can be expected to expand the number of viruses that could be produced as bioweapons using synthetic biology.

The possibility of re-creating known pathogenic bacteria poses a relatively low level of concern, largely because of the high level of technical difficulty. Because they have much larger genomes than viruses, building and booting bacteria would require a great deal of expertise, time, and resources. Given the technical difficulty of this process, actors may find it substantially easier to acquire a pathogenic bacterium through means other than synthesizing them from scratch. (In fact, the same consideration applies to viruses, even if their synthesis is easier than that of bacteria.) In addition, as with viruses, existing mitigation approaches would be expected to be reasonably well equipped to handle an attack using a synthesized known bacterial pathogen. However, two developments could increase the level of concern. If techniques using yeast were to make it far more feasible to boot synthesized bacterial genomes, or if a breakthrough makes it easier to handle large DNA fragments without shearing, the re-creation of bacterial pathogens might warrant increased concern.

The use of synthetic biology to make an existing virus more dangerous poses a medium level of concern. While modifying a virus to change its phenotype may be an attractive option in theory, there are significant barriers to overcome. Such an effort would be working against finely honed virus-host dynamics evolved over millions of years, and a key factor is that modifications to a virus generally lead to attenuation. The barriers are most significant in the Design and Test phases of the Design-Build-Test cycle. While modifying a virus requires significant expertise in viral biology and challenges may be encountered in the Test phase as a result of the inability to ethically test the virus in a human, building the altered virus would be relatively straightforward. High-throughput and directed-evolution approaches could lower the barriers related to the Design phase.

The use of synthetic biology to make an existing bacterium more dangerous poses a relatively high level of concern. This is largely driven by the technical ease of modifying bacterial genomes and the widespread availability of information about the genes involved in traits such as antibiotic resistance and toxin production. Bacteria are routinely modified for a wide variety of beneficial purposes (e.g., to produce biofuels and pharmaceuticals),

TABLE 4-1 Bottlenecks and Barriers That Currently Constrain Capabilities and Developments That Could Reduce These Constraints^a

| Capability | Bottleneck or Barrier | Relevant Developments to Monitor |
|---|---|--|
| Re-creating known pathogenic viruses | Bootling | Demonstrations of bootling viruses with synthesized genomes |
| Re-creating known pathogenic bacteria | DNA synthesis and assembly | Improvements in synthesis and assembly technology for handling larger DNA constructs |
| | Bootling | Demonstrations of bootling bacteria with synthesized genomes |
| Making existing viruses more dangerous | Constraints on viral genome organization | Increased knowledge of viral genome organization and/or demonstration of combinatorial approaches capable of facilitating larger-scale modifications to viral genome |
| | Engineering complex viral traits | Increased knowledge of determinants of complex viral traits, as well as how to engineer pathways to produce them |
| Making existing bacteria more dangerous | Engineering complex bacterial traits | Advances in combinatorial approaches and/or increased knowledge of determinants of complex bacterial traits, as well as how to engineer pathways to produce them |
| Creating new pathogens | Limited knowledge regarding minimal requirements for viability (in both viruses and bacteria) | Increased knowledge of requirements for viability in viruses or bacteria |
| | Constraints on viral genome organization | Increased knowledge of viral genome organization and/or demonstration of combinatorial approaches capable of facilitating larger-scale modifications to viral genome |

^aShading indicates developments that are likely to be propelled by commercial drivers. Some approaches, such as combinatorial approaches and directed evolution, may allow bottlenecks and barriers to be widened or overcome with less explicit knowledge or tools.

and the same techniques and knowledge base would likely prove useful for modifications pursued with a more nefarious intent.

The creation of new pathogens from scratch currently poses a relatively low level of concern, primarily because the knowledge and technologies needed to pursue such an effort are in their infancy. It is likely that a major breakthrough (or more than one) in design capabilities will be required to make this capability a reality.

Relevant developments to monitor for each of these capabilities are summarized in Table 4-1.

5

Assessment of Concerns Related to Production of Chemicals or Biochemicals

Metabolic engineering of microorganisms is a decades-old discipline that has been used to enable manufacturing of a variety of products including fuels, commodity and specialty chemicals, food ingredients, and pharmaceuticals. The core tenets and successes of metabolic engineering are based on the observation that biological systems are inherently chemical systems. A functioning cell, whether of microbial, human, or other origin, is essentially a collection of biochemical reactions taking place within a confined physical space as defined by a cell wall, cytoplasmic membrane, or other enveloping feature. These reactions produce structures that provide both physical form and function. Metabolic engineers have exploited biochemical pathways both to increase the production of compounds an organism naturally produces (e.g., upregulating the production of ethanol by yeast cells) and to coax an organism to produce compounds that are novel to the organism (e.g., rerouting the ergosterol biosynthesis pathway in yeast to produce a plant terpenoid [Kampranis and Makris, 2012]).

Synthetic biology concepts, approaches, and tools have allowed metabolic engineers to pursue an increasingly complex array of chemical products, typically following the overall workflow conceptualized in Figure 5-1. Westfall et al. (2012), for example, engineered yeast to produce artemisinic acid, an antimalarial drug native to the *Artemisia annua* plant. Galanie et al. (2015) added more than 20 genes encoding enzymes nonnative to yeast to the yeast genome in order to produce a variety of plant-based opioids. Microbes have even been engineered to produce compounds for which no naturally occurring biological pathways have been elucidated, such as 1,4-butanediol (Yim et al., 2011), a common industrial chemical also used as a recreational drug.

As the field of synthetic biology endeavors to “improve the process of genetic engineering” (Voigt, 2012), there is a concerted effort across the metabolic engineering community to demonstrate the biological production of increasingly complex molecules while simultaneously developing tools and approaches that reduce the resources required to achieve specific production metrics (e.g., titer, rate, and yield) (NRC, 2015). Hence, it is worth considering how this technology could be misused to produce chemicals or biochemicals for malicious purposes. Such products are likely to fall into one of three categories:

- *Toxins*.¹ Toxins are molecules produced by biological systems that are known to be harmful to humans or other animals. Toxins exhibit wide structural diversity and include small molecules as well as peptides.

¹ The word *biochemical* is used throughout the report to include toxins.

Given that toxins are known to cause harm, they are obvious candidates for engineered synthesis by an actor aiming to do just that.

- *Antimetabolites and small-molecule drugs.* Antimetabolites are compounds that interfere with the normal functioning of cellular metabolism. Although some antimetabolites can be used for therapeutic purposes, as in the use of chemotherapeutic drugs to disrupt metabolic pathways in cancer cells, compounds that target normal functions in healthy tissues can lead to dysfunction or disease. Chemically synthesized small-molecule drugs can also cause dysfunction in healthy tissues. Both antimetabolites and small-molecule drugs may be amenable to synthesis by biological systems.
- *Controlled chemicals.* Synthetic organic chemistry has given rise to a wide variety of chemical compounds with no known biological origin. Many have been essential to advances in human quality of life, whereas others have been used to produce explosives, chemical weapons, and other types of dangerous compounds.

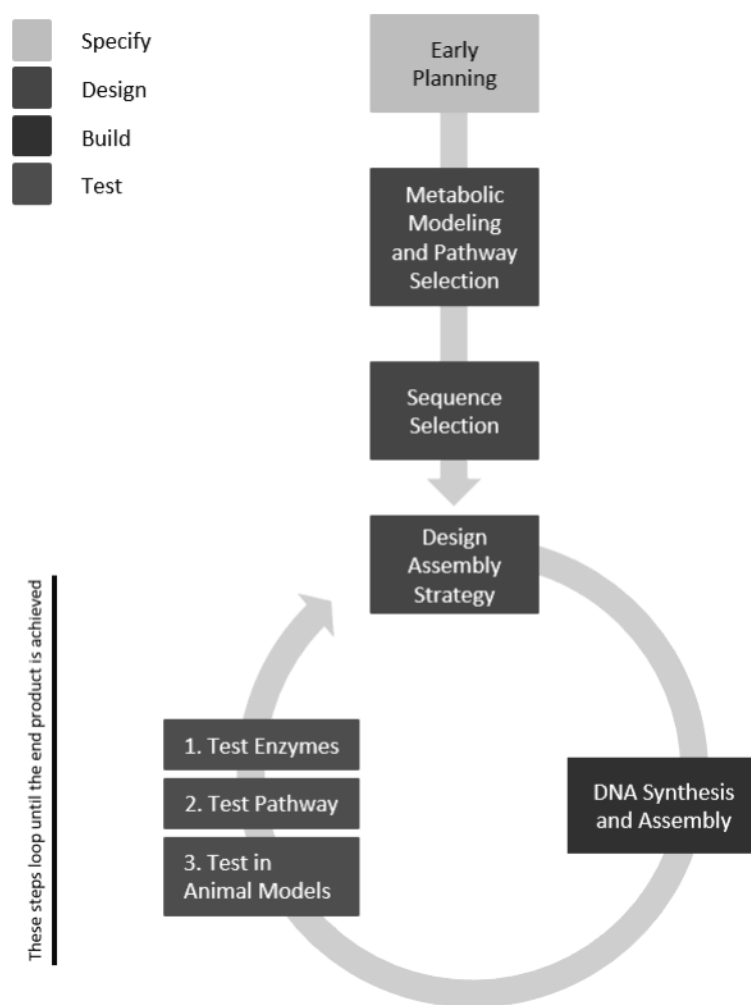


FIGURE 5-1 Activities involved in engineering an organism to produce a desired chemical or biochemical. Considerations in the Design stage may include choice of the host organism, modeling to predict metabolic pathway performance, and bioprospecting for appropriate enzymes to produce the desired product. Multiple rounds of the Design-Build-Test cycle are represented. Testing may first focus on enzyme functionality in early cycles, followed by testing of pathway performance, followed by testing for performance in an animal model in the case of in situ applications.

Some of these compounds (or functionally equivalent analogues) may be accessible through biological synthesis as an alternative to traditional organic chemistry.

While these categories of compounds are instructive in considering end uses, for the purposes of this report it is also useful to differentiate between naturally occurring products (those that are generated in a non-engineered biological host) and manmade products (those that have been chemically synthesized). This distinction affects both the experimental approach and the technical difficulty of using synthetic biology to produce a given target compound. In addition, it is useful to consider the mode of production. For example, target compounds could be produced in small quantities in a laboratory, at large scale in bioreactors (analogous to the industrial production of bio-based chemicals), or even in situ in the human host, such as the production of a toxin by a microbe in the gut microbiome. These various modes offer different challenges with regard to production, delivery, and opportunities for mitigation.

Considering the different types of potential target compounds and the different ways synthetic biology technologies might be exploited to produce them, three main types of activity were identified that are of potential concern: manufacturing chemicals or biochemicals by exploiting natural metabolic pathways, manufacturing chemicals or biochemicals by creating novel metabolic pathways, and making biochemicals via in situ synthesis of target compounds. This chapter assesses the relative level of concern warranted for each of these potential capabilities based on the four framework factors: Usability of the Technology, Usability as a Weapon, Requirements of Actors, and Potential for Mitigation.

MANUFACTURING CHEMICALS OR BIOCHEMICALS BY EXPLOITING NATURAL METABOLIC PATHWAYS

Biochemical compounds naturally produced by plant and microbial cells have been used for centuries as medicinal compounds. These products have been prepared as both plant extracts, in which the active ingredient is one of numerous chemical structures in the formulation, and as high-purity single compounds, made by cultivating the producing organism in large-scale bioreactors and then purifying the output. Such products have been used to treat diseases ranging from microbial infection to hypertension. The opioids, used as analgesics, are now accessible by microbial fermentation, as well, though optimization of the “home-brewing” process has not been rigorously explored (Endy et al., 2015; Galanie et al., 2015).

Each naturally occurring biochemical is the result of a series of chemical reactions that transform simple feedstocks such as glucose into the end products of interest. These transformations are mediated by enzymes encoded by the host organism’s DNA. Because biotechnologies allow the DNA encoding the necessary enzymes to be exploited independent of the original host, it is now possible to make such products without relying on the organism that naturally produces them.

The assessment of concerns related to manufacturing chemicals or biochemicals by exploiting natural metabolic pathways is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for manufacturing chemicals or biochemicals by exploiting natural metabolic pathways | High | High | Medium | Medium-high |

Usability of the Technology (High Concern)

While the production of natural products in microbial hosts is not a trivial endeavor, the core technology required to complete one iteration of the Design-Build-Test cycle for metabolic pathway engineering of a target molecule is readily accessible and relatively easy to use with a basic level of molecular and microbiology expertise. Therefore, the level of concern with regard to this factor is relatively high. Assuming an actor has access to a tractable host organism (e.g., *Escherichia coli*, *Saccharomyces cerevisiae*, *Pseudomonas putida*), the ability to design gene cassettes and insert them into the host, the ability to culture the recombinant host and (as necessary) induce gene expression, and the ability to analyze the resulting products, *attempting* to engineer a metabolic pathway to produce a target toxin or other chemical or biochemical is, on the whole, a relatively straightforward proposition. Although success after one iteration of the Design-Build-Test cycle is probably unlikely, repeated cycles of effort frequently yield improvements in performance.

Of critical importance is whether the pathway, that is, the specific series of chemical reactions leading from a specified starting substrate to the final product, has been fully elucidated. If the pathway is not fully known, this can create a substantial bottleneck or barrier, because a combination of both bioinformatics and experimental techniques would be needed to identify the missing enzymes and reaction steps, necessitating a more advanced level of expertise, more time, and more scientific resources. Difficulty will also increase if a chemical or biochemical is not well tolerated by the host organism engineered to produce the pathway. The difficulty of metabolic engineering also depends on the complexity of the molecule of interest; engineering a pathway to produce structurally simpler molecules will generally be more feasible than engineering a pathway for more complex molecules. For example, the complete biosynthetic pathway for the anticancer drug Taxol remains elusive some five decades after its first discovery in the Pacific yew tree.

Once the pathway is known—and once the genes that encode the pathway enzymes have been specified—the next step is functional expression of the enzymes. This step is often challenging because enzymes transferred from one host to another may lose local structural features that are associated with activity, or they may be separated from essential accessory proteins. The tools of synthetic biology could be used to address these lost structural functions or to provide alternative pathways, but this makes for a more complicated proposition, as discussed below under Manufacturing Chemicals or Biochemicals by Creating Novel Metabolic Pathways. However, if post-translation modifications absent in the new host are essential for enzyme activity, this likely represents an insurmountable hurdle, at least in the near term.

Usability as a Weapon (High Concern)

More than offering new delivery mechanisms or modes of administration, metabolic engineering simply affords access to more material. In short, metabolic engineering in and of itself does not facilitate weaponization, but rather provides a potential means to access larger quantities of harmful material over shorter time frames.

Simply introducing a series of functional enzymes into a suitable host to produce chemicals or biochemicals does not ensure sufficient productivity to warrant concern. Three metrics are essential to assessing the effectiveness of product formation in an engineered organism: productivity (amount of product made per unit of time), titer (concentration of the product external to the engineered organism), and yield (amount of the available feedstock that is converted to product). Whereas such metrics are inconsequential in the native environment (because most biochemicals and peptides are naturally produced in small amounts), these parameters are important to the weaponization of a chemical or biochemical that requires large-scale production. For example, if a toxin is deadly to humans at a concentration of 50 mg/kg, producing that toxin to a titer of 5 mg/L would require someone to ingest at least 10 L of fermentation broth per kilogram of body weight. At a titer of 10 g/L, only 5 mL of broth per kilogram of body weight would need to be ingested. Achieving higher titers allows effective doses to be manufactured in smaller bioreactors, potentially requiring fewer resources. Productivity, titer, and yield determine the volume of cell growth and feedstock needed to make a useful (i.e., harmful) amount of compound, as well as the length of time required for production.

Generally speaking, engineering an organism to increase productivity, titer, and yield becomes progressively

more difficult. At present, engineering microbes to produce toxic small-molecule products in excess of 1 g/L would likely require the dedicated effort of trained metabolic engineers with access to a modern molecular biology laboratory, while a lower titer might be attainable with less expertise and fewer scientific resources. As a result, it can be expected that high-potency toxins would be more desirable targets for malicious actors. However, from the actor's perspective there may also be a trade-off between the relative difficulty of producing a given chemical or biochemical and the amount needed to cause harm. Purity and productivity, as well as the complexity of the target molecule, will also factor into this trade-off. If a compound must have high purity to be effective as a weapon, the difficulty of achieving this level of purity in production or downstream processing (e.g., purifying from lysates) can potentially create a barrier. Low productivity is often related to insufficient substrate concentrations and/or low activity (i.e., the reaction rate is too slow); if enzymatic activity is not sufficiently high to achieve the turnover rates required, even when enzymes are expressed at high levels, additional iterations of the Design-Build-Test cycle may be required to achieve the desired level of productivity.

Once an actor is able to produce a sufficient quantity of a target chemical or biochemical, the predictability of results is likely to be high, assuming the actor has selected a target chemical or biochemical that is already known to cause harm. For example, mass production of botulinum toxin would not require testing of the fermentation product because the effects of its exposure are already known. Indeed, an actor could probably have greater confidence in the effectiveness or lethality of a chemical or biochemical whose pathway is well understood and is produced using synthetic biology as compared to a synthesized pathogen. The latter would definitely require testing to verify that the desired phenotypic results would be achieved.

The scope of casualty expected from a chemical or biochemical compound produced in this way would depend on the amount produced, the potency, and delivery. Chemicals, biochemicals, and toxins do not spread on their own the way pathogens do, and so, effecting a large-scale attack would require delivering a sufficient amount to targeted populations, even if the compound is highly potent. However, there are many potential delivery mechanisms for chemicals or biochemicals, which do not tend to degrade when exposed to the environment the way that pathogens do, and thus would remain potent in a broader array of delivery scenarios than would a pathogen.

In summary, engineering a microorganism to produce a chemical or biochemical by exploiting a natural pathway is considered to pose a relatively high level of concern with regard to usability as a weapon, primarily because of the predictability of the results: Producing a known toxic substance will result in a product with a known toxicity. In addition, chemical or biochemical products are more stable than pathogens. These considerations outweighed the fact that the difficulty of scaling up production to produce large amounts of a substance is a bottleneck or barrier, because there are a number of substances that are highly potent and thus toxic in very small amounts.

Requirements of Actors (Medium Concern)

Generally speaking, the core capabilities for executing a Design-Build-Test cycle in metabolic engineering require a relatively low level of metabolic engineering expertise, especially for a natural metabolic pathway that is already fully elucidated. However, the expertise required depends on the complexity of the pathway and target molecule. Achieving high-level synthesis, especially for difficult targets, does require more expertise and experience; for example, in many cases an actor would need working knowledge of how to knit pathways together into a functioning whole. To fill in the gaps in an incompletely elucidated metabolic pathway, an actor would need access to bioinformatics capabilities in order to analyze genome and transcriptome data, as well as experimental capabilities to detect and identify intermediates. For these reasons, manufacturing chemicals or biochemicals by exploiting natural metabolic pathways is considered to pose a medium level of concern with regard to this factor.

The organizational footprint required depends on the amount of product that is desired (which in turn depends on factors such as potency and titer). Small batches of a chemical or biochemical of interest could be achievable with a relatively small organizational footprint, but scaling up to produce large quantities in a bioreactor would require a larger organizational footprint and more resources.

Potential for Mitigation (Medium-High Concern)

Overall, there is a medium-high level of concern with regard to this factor, primarily driven by the fact that countermeasures are not available for a number of toxins. Lessening the concern slightly is the fact that an attack would be expected to be readily recognized. This assessment assumes that an actor would endeavor to use metabolic engineering to produce compounds with known properties. Because most known biochemicals that could potentially be misused for an attack would naturally be present in very small amounts, the emergence of disease would be a strong indication of purposeful release, thus enabling rapid identification of an attack. However, because the end product would be a chemical or biochemical that is purified away from the organism that produced it, organism-associated signatures would not be available to determine whether the attack resulted from an organism intentionally engineered to produce a dangerous chemical or biochemical, and attributing an engineered organism to a specific actor would be even more difficult.²

The capacity for consequence management depends on the chemical or biochemical used. Governments have developed medical countermeasures to respond to attacks using a subset of known toxins, but there are other toxins that have not been the focus of such efforts. The countermeasures and public health response would be expected to be the same for naturally occurring chemicals or biochemicals and for those created using synthetic biology.

MANUFACTURING CHEMICALS OR BIOCHEMICALS BY
CREATING NOVEL METABOLIC PATHWAYS

While nature has provided a wide array of biochemical compounds that could be exploited for targeted synthesis, enzyme-mediated conversions also can be used to produce chemicals that organisms do not naturally create. Biocatalysis has long been used to produce pharmaceutical intermediates and active ingredients not found in nature (Bornscheuer et al., 2012). It is not always necessary to use living microbial organisms in these processes; instead, purified enzymes can be used in reaction vessels in a manner analogous to traditional organic synthesis. At its core, designing a new biosynthetic pathway involves specifying a series of enzymatic steps that can convert a set starting substrate to the desired end product. In practice, the starting substrate is often a known primary metabolite (e.g., acetyl-CoA) (Savile et al., 2010), and the proposed reaction steps are based on known enzymatic chemistry.

Engineered metabolic pathways that do not follow an existing natural blueprint have been exploited to commercialize biological production of chemical compounds (Yim et al., 2011). The true limits of biological synthesis are unknown, and advances in protein design and engineering are rapidly expanding the repertoire of enzyme-catalyzed reactions (Siegel et al., 2010; Kan et al., 2017). Researchers have also shown that materials typically present in very small amounts in biological systems, such as halogens, can be incorporated into natural products by merging plant and microbial biosynthesis machinery (Rungtaphan et al., 2010). These examples suggest that the range of molecules that may be accessible by biological synthesis is far larger than what has been demonstrated to date.

The assessment of concerns related to manufacturing chemicals or biochemicals by creating novel metabolic pathways is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for manufacturing chemicals or biochemicals by creating novel metabolic pathways | Medium-low | High | Medium-low | Medium-high |

² However, note that the use of isotope ratios for chemical and biochemical attribution has been explored by the Federal Bureau of Investigation (Kreuzer-Martin and Jarman, 2007).

Usability of the Technology (Medium-Low Concern)

Producing a novel metabolic pathway is likely to be significantly more technically challenging than synthesizing a natural metabolic pathway and is likely to require multiple iterations of the Design-Build-Test cycle. Therefore, the level of concern is medium-low with regard to the usability of the technology. The technical challenge stems largely from the fact that engineering novel pathways typically requires engineering enzyme activity, either through rational (computational) design or through directed evolution, to achieve both the activity and specificity required for the pathway of interest. In addition, the enzymes in many cases may be acting on substrates not encountered in nature; in such cases, the likelihood of success is greater if it is structurally similar to the natural substrate of the enzyme being used (Hadadi et al., 2016). For some reactions, it may simply be technologically infeasible to generate high enzymatic activity, but this is likely to be unpredictable, and it may require many Design-Build-Test cycles to determine that one has reached a dead end. Generally speaking, the level of difficulty is likely to be lower if the goal is to engineer a novel pathway that is based on an existing pathway, as opposed to engineering a pathway that is wholly new.

Usability as a Weapon (High Concern)

Considerations related to weaponization, scale-up, predictability of result, delivery, and scope of casualty for novel metabolic pathways are largely similar to those for natural metabolic pathways, and so large-scale production is a barrier or bottleneck. Scaling up production may present additional challenges in the case of novel metabolic pathways if the product is toxic to the cells used to produce it, creating another barrier or bottleneck. In the context of delivery, it may be possible for chemicals created through novel metabolic pathways to be more stable for storage and transport compared to natural biochemicals.

Requirements of Actors (Medium-Low Concern)

While computational tools and established methodologies exist for creating new metabolic pathways, metabolic engineering is still largely an “art” rather than a “science.” Because intuition continues to play a significant role in the successful execution of experimental designs, creating functional novel metabolic pathways is likely to require a higher level of expertise and experience than exploiting natural pathways would. In particular, if a novel pathway requires enzymes to act on novel substrates, expertise in protein engineering (which is beyond the typical skill set of an experienced metabolic engineer) would also be required. Both the knowledge about how to design novel pathways and knowledge of how to engineer enzyme activity are bottlenecks or barriers in this space. Therefore, the level of concern with regard to this factor is medium-low.

Potential for Mitigation (Medium-High Concern)

Considerations related to mitigation capabilities for chemicals or biochemicals manufactured by creating novel metabolic pathways are largely similar to those for chemicals or biochemicals created through natural metabolic pathways.

MAKING BIOCHEMICALS VIA IN SITU SYNTHESIS

The human microbiome, particularly the gut microbiome, has been a target for metabolic engineering. Gut microbes influence the metabolism of their host and are capable of producing a wide variety of biochemicals. While the extent of the influence of the microbiome on host metabolism remains an active research area, there has already been significant progress toward engineering gut microbes for therapeutic purposes. Engineered microbes are currently being prepared for clinical trials for the treatment of metabolic disorders (Synlogic, 2017), although engineering high flux through a metabolic pathway remains undemonstrated.

As this research gains steam, it is worth considering whether the human microbiota could be exploited to

make biochemicals (within the cells of commensal organisms) and deliver them to human hosts to cause harm. In addition to the gut microbiome, the skin microbiome could be another potential avenue for in situ synthesis of such compounds. Related concepts include the manipulation of the human microbiome to cause dysbioses or as an avenue for horizontal gene transfer (see Chapter 6, Modifying the Human Microbiome). Environmental dispersion of a microorganism capable of producing toxins, antimetabolites, or controlled chemicals may also be considered a potential in situ delivery mechanism, one whose outcome would be difficult to predict. The basic principles of pathway engineering in a microbe are the same whether the intention is to culture the organisms in large vessels followed by purification of the molecules of interest or to introduce the organisms into the environment or a human host for in situ production and release of a biochemical. However, the scope of the engineering effort can vary substantially since manufacturing in vessels is likely to require that much higher production titers be achieved. For example, nanograms of a sufficiently toxic material delivered in situ could be sufficient to produce a harmful effect compared to tens of grams per liter needed for cultivation in and purification from fermentation vessels. This difference is important to consider in assessing concerns.

The assessment of concerns related to making biochemicals via in situ synthesis is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for making chemicals or biochemicals via in situ synthesis | Medium-high | Medium | Medium | High |

Usability of the Technology (Medium-High Concern)

From an engineering perspective, creating a microbe capable of in situ biological synthesis of a biochemical presents many of the same opportunities and challenges as engineering metabolic pathways for the production of chemicals or biochemicals in a bioreactor, though there are some additional challenges, as well. While productivity, titer, and yield can typically be measured in the process of manufacturing a chemical or biochemical product in a bioreactor, conditions in the microbiome, for example, are quite different from those present in the laboratory. This makes it difficult to predict and control whether productivity, titer, and yield measurements in the laboratory will translate to similar numbers once the microbe is delivered to the microbiome (or environment). Many Design-Build-Test cycles, including a substantial amount of testing in both cell cultures and in animal models, are currently needed to obtain engineered gut microbes with functional gene circuits (Lu et al., 2009; Kotula et al., 2014; Mimeo et al., 2015; Matheson, 2016). One potential way to expedite development and reduce the need for multiple rounds of resource-intensive in vitro and in vivo testing would be to expose human subjects to large libraries of prototype microbes, then sequence the microbiome content to identify the successful prototype microbes if toxicity is observed. However, this library approach has important limitations. For example, a prototype microbe capable of producing high titers of a toxin if introduced to the gut as a monoculture could be effectively diluted by the presence of large numbers of ineffective prototype microbes, making it difficult to detect and identify the successful prototype microbe. In addition, it is possible that a microbe that produces high titers of a toxin would grow more slowly than prototype microbes that produce little or no toxin, making it difficult to separate signal from noise. Finally, the current state of the art in gut microbiome sequencing and assembly does not guarantee that a successful prototype strain could be correctly constructed and differentiated from all other introduced library strains. Nonetheless, the fact that many organisms harbor their own toxins as part of their infective life cycle means that it should not be impossible to align pathogenicity and evolutionary fitness, and indeed one of the easiest means of establishing a toxin in situ may be via an already known pathogen, as discussed under Usability as a Weapon, below, and in Chapter 4, Box 4-2.

Overall, the knowledge needed to manipulate organisms in the gut and skin microbiome remains limited, as further discussed in Chapter 6, Modifying the Human Microbiome, and it is possible that unforeseen challenges in

producing biochemicals in situ will emerge in the coming years. However, the field has been advancing quickly. Already, researchers have demonstrated the ability to manipulate some human gut microbes, and the use of the microbiome for delivery of pharmaceuticals is an active area of research. Thus, the high rate of development and investment in this field leads to a medium-high level of concern with regard to this factor. It will be important to monitor for research breakthroughs that exacerbate opportunities for misuse in this area, as well as breakthroughs in understanding.

Usability as a Weapon (Medium Concern)

Usability as a weapon is considered of medium concern, largely due to current limitations in the ability to make introduced microbes persist in the microbiome. However, microbiome engineering is an active area of research, and significant advances, such as a demonstrated ability to cause persistent changes in the gut microflora, would cause the level of concern to rise.

The gut microbiome is known to host thousands of gene clusters, and products of these clusters have been shown to be present in the gut at high micromolar concentrations (Donia and Fischbach, 2015). Therefore, it should be possible to engineer gut microbes to produce harmful small molecules at similar levels. However, despite the presence of these natural pathways in the microbiome, the principles behind engineering similar pathways to produce other products in situ have not been determined. Engineering the production of a toxin with sufficient titer, produced over a long enough time to be harmful to the host, is not necessarily straightforward. Furthermore, after being delivered into the host microbiome, the engineered microbe would need to colonize and persist to have a long-term effect. Experiments with attenuated vaccine strains suggest that it is necessary to eliminate some existing microbes in order to allow an introduced microbe to persist in the gut, adding to the complexity of purposefully infiltrating a host microbiome. A perhaps more likely scenario is that existing gut or skin microbes could be manipulated to increase their natural production of a harmful compound or to resist antibiotics or other countermeasures, thus allowing delivery of an agent without the barrier of infiltrating the native microbiome with a new microbial species. In addition, it is possible that a pathway lodged on a broad-host-range vector might be horizontally transferred to native species following transient introduction on a microbe that was otherwise unlikely to colonize; the horizontal transfer of in situ engineered pathways is further considered in Chapter 6, *Modifying the Human Microbiome*.

Although the chemical product would be manufactured by cells, bioreactors or flasks would likely be required to produce a sufficient number of cells to enable delivery to the target human population. Microbes engineered to secrete highly potent biochemicals, which could cause greater damage in smaller quantities, would warrant greater concern than those engineered to produce lower-potency chemicals. But effectively delivering engineered microbes to the human target would still present significant barriers. Cold War-era studies on the weaponization of bacteria remain relevant to this concept. Contamination of food could be an efficient method of dispersal, but could be thwarted by standard food safety measures such as cold storage, cooking, and mechanisms to limit the spread of contaminated food. The scope of casualty from in situ biosynthesis would be expected to be relatively low, because the agent would need to be delivered to each individual and then persist in the gut or skin long enough to cause harm. That said, the ability to slightly or gradually modify human physiology and behavior via even low-level production of compounds could be extremely debilitating to a modern nation-state.

Requirements of Actors (Medium Concern)

Engineering microbes to actively secrete products in the microbiome would generally require a higher level of expertise than engineering a natural metabolic pathway but less sophistication than designing a novel metabolic pathway, leading to a medium level of concern with regard to this factor. Because multiple iterations of the Design-Build-Test cycle would be needed, actors would likely require access to significant laboratory resources over a long period of time. On the other hand, in situ synthesis presents fewer barriers with regard to scale-up and downstream processing than the production of chemicals or biochemicals in a bioreactor, and once a sufficient

engineered microbe is developed, producing and delivering a small quantity would not require a great deal of technical expertise.

Potential for Mitigation (High Concern)

The challenges of attribution and the difficulty of identifying and stopping an attack based on in situ synthesis of biochemicals lead to a relatively high level of concern with regard to this factor. Policies and procedures related to the containment of natural foodborne pathogen outbreaks should transfer well to the containment of engineered toxin-producing gut microbes. Indeed, the presence of strong public health infrastructure for food safety and response to contaminated-food outbreaks may deter skilled actors from pursuing an attack with engineered gut bacteria in favor of other attack vectors. In addition, while engineering microbes to resist traditional countermeasures (such as the use of broad-spectrum antibiotics) could increase the casualty rate, containment and isolation of contaminated facilities would be expected to limit the spread of such agents. However, the delivery of engineered microbes to the gut via food is not the only potential attack vector or means of delivery. The development of an engineered microbe that could infiltrate the skin microbiome, or the development of a high-efficiency method of delivering gut microbes, could be less vulnerable to existing mitigation measures and thus significantly increase the level of concern warranted. However, these delivery modes are currently theoretical.

Regardless of the effectiveness of public health infrastructure for containing an attack, it could be extremely difficult to recognize an attack—that is, to differentiate between a natural disease outbreak and an intentional introduction of engineered microbes into the microbiomes of affected people. This difficulty is the primary driver of the relatively high level of concern related to the potential for mitigation. Some types of attack would be easier to recognize than others; for example, the presence of an unlikely gut toxin or extremely high resistance to available countermeasures may be more easily recognized as signs of an attack, while tracing an effect that is not a classical gut problem (e.g., opioids made in the gut) to engineered gut microbiota would be a substantial task.

In contrast to the other applications of metabolic engineering discussed in this chapter, the genetic material of the engineered microbe would in the case of in situ synthesis remain present in the weaponized product. Sequencing clinical samples of impacted individuals could allow investigators to identify the genetic sequences or organisms used in an attack. However, such an effort would face significant technical challenges. First, if the engineered microbe is present in low abundance, most of the sequence data in a sample would come from non-engineered commensal microbes. Compounding this, only a small amount of the genome of an engineered microbe would be expected to contain new DNA. For example, an engineered *Escherichia coli* genome could contain fewer than 10 heterologous genes, which would need to be detected within the rest of genome, which contains more than 4,000 genes. The high complexity and variability of the gut microbiome composition increases the potential that uncharacterized genes present in the sequencing data could be confused with transgenes.

Even if the sequence of an engineered pathway could be identified in a clinical sample, it may still be difficult to trace the attack to the actors responsible. One potential approach would be to attempt to identify the vendor that produced the synthesized DNA. However, with DNA synthesis technology becoming increasingly accessible, it may become difficult to query all companies capable of producing synthetic DNA. Furthermore, assembly of synthetic DNA from nucleotides could obviate the need for DNA synthesis from a commercial provider. While investigative work in tracing the engineered microbes to their source is likely to be more informative than focusing on the transgenic DNA sequences, the sequences would be extremely important to connecting suspected actors to the weapon material, if matching materials in the actor's laboratory were available.

SUMMARY

- Synthetic biology enables new ways to create harmful chemicals and biochemicals, including toxins.
- Chemicals and toxins produced via manipulation of biological components may be high potency, requiring small amounts to cause harm, or low potency, requiring larger amounts. Although synthetic biology can facilitate development in either case, high-potency chemicals or biochemicals require less downstream expertise with regard to production and delivery. Producing and delivering sufficient amounts of lower-potency chemicals or biochemicals would require greater expertise and more advanced technology to achieve both suitable strain performance metrics and production at appropriate volumetric scales.
- The production of chemicals or biochemicals that do not occur naturally (and do not have a published known metabolic pathway) requires specific expertise due to the challenges associated with enzyme engineering and elucidating and specifying metabolic pathways.
- In situ production of biochemicals is of higher concern, largely due to limited mitigation capabilities for such a novel approach, including a limited ability to recognize an attack and a potential lack of effective countermeasures.

This chapter considers various ways in which synthetic biology technologies could potentially be applied to produce chemicals and biochemicals such as toxins, antimetabolites, small-molecule drugs, or controlled chemicals for use in an attack. Broadly, the use of microbes to synthesize agents in situ presents the greatest level of concern, the synthesis of agents using naturally occurring metabolic pathways warrants a medium to high relative level of concern, and the engineering of novel metabolic pathways poses a medium level of concern.

It will be important to continue to monitor developments in the manipulation of the human microbiome because efforts in the pharmaceutical arena are likely to propel advances and reduce bottlenecks and barriers as the field continues to progress (see Table 5-1). Although the level of certainty around the in situ manufacture of biochemicals via the gut or skin microbiome is lower than the level of certainty involved in the other metabolic engineering processes described in this chapter, manipulation of the microbiome is an active and quickly advancing area of research. Overall, this potential capability warrants a higher level of concern, because an attack effected through manipulation of the human microbiome could be difficult to recognize and trace. However, understanding of microbiome dynamics is still relatively limited, and it would likely take a relatively high level of expertise and many iterations of the Design-Build-Test cycle to develop a microbe capable of colonizing the human host microbiome, manufacturing the biochemical in sufficient quantities, and persisting long enough to cause harm.

The primary drivers of the medium to high relative level of concern for the potential exploitation of naturally occurring metabolic pathways are the relatively high level of knowledge available, the relatively low level of technical expertise required, the availability of multiple delivery mechanisms, and the difficulty of tracing the source of an attack. Exploitation of naturally occurring pathways could be an option for attackers because it is easier, in general, to use microbes to manufacture complex chemicals or biochemicals than to use chemical synthesis techniques. However, scalability remains a bottleneck, and manufacturing large enough quantities of the chemical or biochemical to effect a large-scale attack would require a large organizational footprint. Given this, a more likely application of this approach may be to manufacture drugs, such as opioids. The difficulty of this approach also depends heavily on the complexity of the chemical or biochemical of interest and of the metabolic pathway for producing it. For some target chemicals or biochemicals, an actor may conclude that cultivating the native host organism may be more feasible than using metabolic engineering to produce a biochemical in a bioreactor (e.g., cultivating *Clostridium botulinum* instead of heterologous production of botulinum toxin).

The development of novel metabolic pathways to produce chemicals is a technically challenging proposition that would require expertise in both metabolic engineering and protein engineering in order to develop the necessary enzymatic activities, and further efforts to make the novel pathway yield a sufficient amount of product for

TABLE 5-1 Bottlenecks and Barriers That Currently Constrain the Capabilities Considered and Developments That Could Reduce These Constraints^a

| Capability | Bottleneck or Barrier | Relevant Developments to Monitor |
|--|--|--|
| Manufacturing chemicals or natural metabolic pathways | Tolerability of toxins to the host organism synthesizing the toxin | Pathway elucidation, improvements in circuit design, and improvements in host (“chassis”) engineering to make toxins tolerable to the host organism synthesizing the toxin |
| | Pathway not known | Pathway elucidation and/or demonstrations of combinatorial approaches |
| | Challenges to large-scale production | Improvements in intracellular and industrial productivity |
| Manufacturing chemicals or biochemicals by creating novel metabolic pathways | Tolerability of toxins to the host organism synthesizing the toxin | Pathway elucidation and/or improvements in circuit design and/or improvements in host (“chassis”) engineering to make toxins tolerable to the host organism synthesizing the toxin |
| | Engineering enzyme activity | Increased knowledge of how to modify enzymatic functions to make specific products |
| | Limited knowledge of requirements for designing novel pathways | Improvements in directed evolution and/or increased knowledge of how to build pathways from disparate organisms |
| | Challenges to large-scale production | Improvements in intracellular and industrial productivity |
| Making biochemicals via in situ synthesis | Limited understanding of microbiome | Improvements in knowledge related to microbiome colonization of host, in situ horizontal transfer of genetic elements, and other relationships between microbiome organisms and host processes |

^aShading indicates developments that are likely to be propelled by commercial drivers. Some approaches, such as combinatorial approaches and directed evolution, may allow bottlenecks and barriers to be widened or overcome with less explicit knowledge or tools.

an attack. Multiple iterations of the Design-Build-Test cycle would be required. The difficulty would be reduced if the novel metabolic pathway were to use steps, enzymes, or substrates from a naturally occurring pathway, and indeed, recent advances in protein design and engineering have rapidly expanded capabilities for engineering novel metabolic pathways. The most feasible metabolic routes will be those that have been already demonstrated elsewhere (e.g., in the academic literature), because recapitulating an engineered pathway is substantially more tractable than developing a pathway from scratch. However, even where biological synthesis is feasible for producing controlled chemicals or other products, traditional chemical synthesis may prove to be a more reliable, cost-effective, and surreptitious means to do so when the involved pathways are novel. An actor skilled in the art of metabolic engineering who is capable of engineering high-titer strains and has access to the right scientific resources is expected also to be sufficiently skilled to access, and potentially opt for, these other options.

Relevant developments to monitor for each of these capabilities are summarized in Table 5-1.

6

Assessment of Concerns Related to Bioweapons that Alter the Human Host

While we typically think about biodefense in terms of either pathogens (Chapter 4) or biochemicals (Chapter 5), technological advances are now making possible additional capabilities and means of attack that are more closely related to the human body itself. The study included consideration of how increased knowledge about the microbiome and immune system may enable new means of delivering an agent; the potential for incursions into the human host through means not typical of pathogens or toxin-based bioweapons, such as through genetic modification; and how genes themselves may potentially be used as weapons. While some of these potential activities overlap with the activities discussed in previous chapters, it is valuable to consider them from a host-centric angle to assess how advances in knowledge and biotechnology tools might further alter the landscape of vulnerabilities and weapons available for exploitation by malicious actors.

MODIFYING THE HUMAN MICROBIOME

Human health is highly dependent upon the human microbiome—the microorganisms that live on and within us, especially those associated with the gut, oral cavity, nasopharyngeal space, and skin. These populations of microbes are likely far easier to manipulate than the human host itself, making the microbiome a potentially accessible vector for attack. The human microbiome is the focus of a great deal of academic and commercial research, and microbiome manipulation is an area that is rapidly developing, as also discussed in Chapter 5. Several possible ways the microbiome could be manipulated to cause harm were considered; these possibilities were analyzed, in aggregate, to determine the level of concern warranted.

Delivery of harmful cargo via the microbiome. As discussed in Chapter 5, the engineering of microorganisms to produce hazardous chemicals or biochemicals (including toxins) poses a medium to high level of concern and the potential for making chemicals or biochemicals in situ via the microbiome warrants a high level of concern. The microbiome could be used as a vector for other types of harmful cargoes, as well. For example, microbes could be modified to produce functional small RNAs (e.g., microRNAs [miRNAs]) that could be transferred to the host via the gut or skin microbiome¹ to cause a variety of health impacts.² Microbes also could potentially be engineered to horizontally transfer a genetic cargo to the native microbiome to, for example, cause a host's

¹ The transfer of small RNAs has been demonstrated in other organisms (Zhang et al., 2012), and small RNAs and other nucleic acids derived directly from the diet have been found circulating in higher organisms (Yang et al., 2015).

² In human skin, application of anti-tyrosinase siRNAs leads to temporary changes in skin pigmentation (Kim et al., 2012).

own well-established microbes to produce a harmful biochemical. In such a scenario the harmful agent would be manufactured by organisms in the established microbiome, so the engineered microbe would need to infiltrate and persist within the microbiome only long enough to transfer its cargo to a sufficient number of native microbes. Thus, this approach would circumvent the challenges associated with establishing engineered microbes in otherwise occupied niches. There are many known instances of natural horizontal transfer events that result in the production of toxins (Kaper et al., 2004; Strauch et al., 2008; Khalil et al., 2016). It may be possible to harm a population by enhancing the spread of vectors or phage (viruses targeting bacteria [Krishnamurthy et al., 2016]) carrying such genetic cargoes. Synthetic biology methods could advance such a capability, for example, through the engineering of toxin:antitoxin couples that would help ensure retention of plasmids. It is also conceivable that microbes could one day be engineered to horizontally transfer genes directly to human cells.

Use of the microbiome to increase the impact of an attack. The microbiome can also potentially be exploited to design a more effective bioweapon or increase the impact of an attack. Knowledge of the human microbiome could be used to modify pathogens or their delivery mechanisms to allow more efficient propagation within or between populations, for example, by taking advantage of the frequent exchange of bacteria between humans and animals. In particular, domestic animals could be used as carriers for engineered agents transmitted via the microbiome. For example, engineered dog or cat microbiomes could be established via adulterated feedstocks or via purposeful contamination of populations in animal shelters or pet stores and then subsequently transmitted to humans. Natural transfers resulting from animal-human contact, such as the transfer of the parasite *Toxoplasma gondii* from cats to humans and the transfer of *Campylobacter* from dogs to humans, illustrate the feasibility of this approach (Jochem, 2017). Similarly, research into the role of the microbiome in pathogenesis could provide a roadmap as to how to generate improved pathogens that are better supported by their microbial peers. Studies involving wide-ranging transposon- or CRISPR-based deletion libraries of pathogens (Barquist et al., 2013) have provided many insights into pathogenesis that might have dual-use implications, and such libraries could prove useful in identifying which genes productively or specifically interact with endogenous flora to better establish a pathogen.

In addition to using the microbiome to spread toxins and pathogens, manipulating the microbiome might also prove to be a useful adjunct for other biological threats. Recent research shows, for example, that eukaryotic viruses utilize bacteria to improve their chances of infection (Kuss et al., 2011). It is also conceivable that an actor could introduce an initial agent into a population in order to trigger widespread treatment with broad-spectrum antibiotics and then take advantage of the treated population's "clean slate" to introduce or expand an engineered organism via the (now disrupted) microbiome. An actor taking this two-step approach could even incorporate antibiotic or antiviral resistance elements into the initial attack.

Engineered dysbiosis. Our ever-increasing understanding of the human microbiome may lead to opportunities for engineered dysbiosis—that is, the purposeful perturbation of the normally healthy microbiome. This could be accomplished either by causing a known dysbiosis or engineering a new one, and in either case would likely involve introducing otherwise nonpathogenic microorganisms that then lead to diminutions in human health and performance. Since the microbiome likely plays a key role in human immunity (Kau et al., 2011), dysbioses could also potentially be used to cause longer-term debilitation of a population's ability to defend against disease. Gut, oral, nasal, and skin microbiomes could be targets for such an approach. The degradation of military readiness due to continued operations in harsh climates is an ongoing issue. This situation could be made much worse by targeted additions to or alterations of the skin microbiome that lead to heightened chafing, rashes, windburn, and itchiness. While these are seemingly minor concerns, over time they could degrade military capabilities to the point of impacting readiness.

The assessment of concerns related to modifying the human microbiome is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for modifying the human microbiome | Medium-low | Medium | Medium | Medium-high |

Usability of the Technology (Medium-Low Concern)

Engineering the microbiome for any of the purposes described above would be difficult in the near term, leading to a medium-low level of concern with regard to this factor. Given the current level of understanding of the microbiome, the genetic modification(s) required to effect desired phenotypic changes are not yet certain. Achieving desired phenotypic results might require the introduction of particular bacterial species or strains and/or particular genetic modifications of these species or strains. In most cases, microbiome engineering is likely to be further complicated by the need to make multiple genetic introductions or edits involving multiple symbiotic microbiome species. Activities in this area may also be hampered by limited understanding of the genomic diversity and plasticity of microbial communities. Today's genomic databases are built around consensus sequences and do not adequately store or link genomic variations from a single sample. The surprisingly large differences in genomic plasticity observed when the U.S. Food and Drug Administration first applied whole-genome sequencing to trace an *Escherichia coli* outbreak underscore the inadequacy of this approach (Eppinger et al., 2011) and also suggest the difficulties inherent in engineering the microbiome.

There are similar barriers to understanding how to rationally manipulate the environment to encourage particular microbial compositions. For example, the vast differences in human diets worldwide create a plethora of different microbial environments that would be difficult to uniformly engineer. Even if insertion of a pathogenic microbe were possible, metabolism in culture is so different from metabolism in a host that if a given metabolic pathway was altered to achieve a particular phenotype, alternative or secondary pathways might be uniquely turned on in the context of a human host, thus potentially damping or thwarting the desired microbiome phenotypic engineering outcome. However, the microbiome is an extremely active area of research, and capabilities are advancing rapidly, particularly with regard to understanding how environmental perturbations affect species representations (Candela et al., 2012; Ghaisas et al., 2016) and with regard to the development of phages to target bacteria. It will be important to monitor new developments as the enormous interest in the impact of human commensals on human health continues to drive research and investment and will impact the current bottleneck of limited microbiome understanding.

Usability as a Weapon (Medium Concern)

There are many known routes for the introduction of bacteria into populations; the gut, mouth, nasal, or skin microbiomes could potentially be infiltrated through ingestion, dermal, or other exposure routes via a wide variety of avenues, from contaminated food or water to airborne sprays. For the warfighter, the uniformity of the food supply chain may make food of particular concern as a vector for attack; additionally, products such as probiotics and herbal supplements, routinely used by many warfighters (Hughes et al., 2010; Daigle et al., 2015) could be exploited. It also may be possible to engineer a bioweapon to target populations with a specific microbiome profile; any adversary that begins to better parse, store, and analyze the data that are increasingly being collected about human microbiomes will also be in a better position for probabilistic targeting of microbiome threats (see also Chapter 7, Targeting). However, the predictability of the results for manipulation of the microbiome will be low and, unlike conventional pathogens, the opportunities for dissemination via human-to-human transmission are reduced. On balance, the availability of routes to introduce bacteria tempered by the lack of predictability leads to an overall level of medium concern for this factor.

Requirements of Actors (Medium Concern)

The probiotics industry is well established and highly distributed; probiotics are being engineered and manufactured by people around the world with relatively low levels of scientific expertise at small-scale facilities using basic equipment. Once a successful microbiome engineering approach is established, subsequent production of bioweapons could likely be achieved with a relatively small organizational footprint. However, a high level of expertise would likely be needed to perform the engineering required. On balance, the expertise required to overcome the technical challenges in combination with the low organizational footprint leads to a medium level of concern for this factor.

Potential for Mitigation (Medium-High Concern)

The ability to recognize and respond effectively to an attack involving the microbiome would likely vary depending on the approach used. Given the still nascent understanding of the succession of microbial populations, the targeted manipulation of the human microbiome is, generally speaking, likely to be difficult to detect or attribute. The effects of an engineered threat, stealthily introduced, might be easily passed off as part of a normal change in microbial composition, particularly if the effects are slow acting or chronic phenotypes (e.g., mental health deficits, immune suppression, skin rashes). If an attack were detected, the individuality and plasticity of the human microbiome would likely make attribution difficult. Additionally, given the proliferation of facilities involved in manufacturing probiotics, it could be difficult to distinguish intentional production of harmful probiotics from natural issues arising from contamination or other breakdowns in normal production quality control. However, the gut and other microbiomes are robust and regularly reestablish microbial equilibria after perturbation, and existing antibiotics may well be an effective countermeasure against engineered microbes. As a result, treating attack victims could be relatively straightforward, and existing public health and outbreak response measures could be well positioned to contain an attack. While the introduction of antibiotic resistance genes might restrict the possibilities for treatment, this problem differs little from the traditional concerns over the spread of antibiotic resistance in populations and can potentially be overcome through the use of novel antimicrobials, especially in small cohorts. The overall level of concern for this factor is medium-high; the high level of concern that such an attack would be difficult to detect is reduced somewhat by the ability to treat if it were detected.

MODIFYING THE HUMAN IMMUNE SYSTEM

Human immunity is the bulwark for protection against infectious disease. Two basic systems respond to the vast array of threats in the natural environment. The first is the innate immune system, a collection of nonspecific protective mechanisms triggered by pathogen-associated molecular patterns, such as lipoteichoic acid from Gram-positive bacteria or unmethylated CpG sequences in viral DNA. The second is the adaptive immune system, which generates highly specific antibody and T-cell responses tailored to individual diseases and disease variants. Many natural pathogens manipulate the human immune system, both by suppressing the immune response (e.g., immunodeficiency viruses) and by upregulating certain responses (e.g., respiratory syncytial virus, which induces the immune system to favor a response involving Type 2 T helper cells [Th2] and subsequently increases the proclivity toward asthma [Lotz and Peebles, 2012]). These examples suggest that it may be feasible to develop a bioweapon capable of manipulating or “engineering” the immune response. Several potential forms for such a bioweapon were considered:

Engineering immunodeficiency. Manipulating a target population to have decreased immunity could increase the impact of a biological attack. This goal could be pursued either by manipulating a pathogen to simultaneously reduce immunity and cause disease (Jackson et al., 2001) or by separately introducing an immune-suppressing agent and a bioweapon into a target population. Agents used to cause immunodeficiency could be pathogens (e.g., the insidious spread of HIV [human immunodeficiency virus]) or chemicals (see NRC [1992] and IPCS [1996] for discussions of chemicals that contribute to immunotoxicity). It is also possible that a disease agent could be tailored to the immune state of a population, either by engineering the agent to avoid extant adaptive or innate

immune barriers or by actually taking advantage of those barriers (for further discussion see Chapter 7, Health-Associated Data and Bioinformatics).

Engineering hyperreactivity. The flip side of engineering immune deficiencies would be to attempt to cause immune hyperreactivity. Both pathogens and chemicals have been demonstrated to create a cytokine storm, a dangerous state that results from a positive feedback loop in the immune response. It may be possible to engineer an agent to purposefully trigger such a cascade. For example, some have suggested that the introduction of anthrax lethal toxin into a more benign disease vector could trigger a cytokine storm (Muehlbauer et al., 2007; Brojatsch et al., 2014; however, see Guichard et al., 2012 for a differing point of view). Similarly, the fact that there are already widespread responses in the human population to a limited number of well-known allergens (ACAAI, 2017) may provide a means of engineering biological threats that would trigger life-threatening IgE-mediated immune responses. The development and testing of new immunotherapies could also provide a roadmap for potentially engineering threats; for example, actors could learn from clinical studies in which anti-CD28 antibodies caused life-threatening cytokine storms (Suntharalingam et al., 2006).

Engineering autoimmunity. Natural autoimmune diseases cause significant disability and death. It may be possible to engineer a disease that causes the body to turn on itself. Mouse models for the stimulation of autoimmunity now exist. For example, Experimental Autoimmune Encephalomyelitis, which mimics the symptoms of the human malady multiple sclerosis, has been induced in mice by immunization with antigens that cause an immune response (autoantigens; see Miller et al., 2007). Normally, such self-immunization is prevented by the mechanisms that ensure exclusion of antibodies and T-cells that are self-reactive, but some pathogens may present antigens that are similar enough to the body’s own proteins that the original immune response spreads from the pathogen to the new human target. Research into checkpoint inhibitors, compounds designed to unleash the human immune system to eradicate tumors, could also potentially inform efforts to purposely engineer autoimmunity. By overstimulating the immune system, checkpoint inhibitors have been shown to lead to autoimmunity, often in the form of colitis (June et al., 2017). In addition, particular compounds have been shown to lead to an autoimmune disease of the liver (Tanaka et al., 2017, 2018). One potential route of attack could be to introduce such compounds via the microbiome.

The assessment of concerns related to immunomodulation is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|--|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for modifying the human immune system | Medium | Medium-low | Low | High |

Usability of the Technology (Medium Concern)

It is difficult to predict precisely the impact of engineering on a system as complex as the immune system. We are only now beginning to more fully understand the mechanisms for how the immune system recognizes foreign antigens, and many immune mechanisms, such as how immune memory guides future responses, remain opaque. In addition, much of the research in this area is on animals, and the results do not necessarily map well to humans. Furthermore, while there has been an explosion of new research into the causes of autoimmunity, the onset of autoimmune disease remains idiosyncratic (Rosen and Casciola-Rosen, 2016), and it would likely be difficult to create immunomodulatory weapons capable of causing reliable effects in populations as genetically and immunologically diverse as the United States. In particular, while an immune deficiency virus pandemic has emerged naturally, engineering the spread of immune deficiency is currently difficult to imagine.

However, even undirected efforts in this area could be successful enough to warrant concern. In experiments in which mousepox was augmented with interleukin-4 (IL-4) (Jackson et al., 2001), earlier studies had already discerned that vaccinia virus altered with IL-4 increased virulence in mice (van den Broek et al., 2000), but it came

as a surprise that the altered mousepox virus could also overcome vaccination against mousepox. The failed clinical trial of anti-CD28 antibodies, in which patients suffered life-threatening cytokine storms after receiving doses 500 times lower than those shown safe in mouse models (Suntharalingam et al., 2006), offers another example. Although modeling studies indicated that the doses used would nearly saturate the T-cell population of a human (suggesting the potential for overactivation), the dramatic outcomes highlight the potential for inadvertent immune hyperreactivity as well as the dual-use potential of immunomodulation research. The concept of engineering a cytokine storm, especially in susceptible subpopulations, may become a concern when coupled with increasing knowledge of the immune system. For example, the growing knowledge of superantigens that hyperstimulate immunity could further increase the feasibility of such activities.

Our understanding of human immunity also represents an increasing, but unknown, area of concern. For example, with the advent of next-generation sequencing, the range of both B-cell and T-cell responses to vaccines can now be described in molecular detail. Similarly, the effectors of the pattern recognition receptors of the innate immune system are being defined to the point that engineering responses, both therapeutic and otherwise, are possible (Brubaker et al., 2015; Macho and Zipfel, 2015). In addition, the continuing explosion of work in immunotherapy broadly could potentially create a roadmap for the development of immunomodulatory weapons. As understanding of this phenomenon improves and as the ability to engineer protein structures improves, the opportunities for creating synthetic simulacrum of antigens already known to be present in autoimmune diseases will increase. The opportunities to engineer autoimmunity are likely tempered by the diversity of potential autoantigens that can be exploited, although this could also be viewed as a means of disease targeting as more and more personalized health data become available (see Chapter 7, Health-Associated Data and Bioinformatics).

On balance, given the challenges and both near- and longer-term opportunities, there is a medium level of concern with regard to usability of the technology for the variety of ways in which immunomodulation might be employed as a bioweapon.

Usability as a Weapon (Medium-Low Concern)

The connections between factors capable of influencing immunity and the actual immune response of individuals remain poorly understood. Although it is possible to imagine generic degradations to, or overstimulation or mis-stimulation of, the human immune system, it will initially be very difficult to target such threats to particular individuals or populations, and thereby to have a clear and predictable path to an overall impact on a population's health or on military readiness and response. However, although immunomodulation might not necessarily be the most effective approach for an adversary seeking to effect large-scale and immediate death or debilitation, this approach could nonetheless undermine a nation's capabilities. The 1918 influenza pandemic, likely abetted by an interplay between viral infectivity and poor public health, was a major factor in military preparations for the first World War (Byerly, 2010); this historical example serves as a reminder that a general decrease in immunity would even today have strategic consequences for the military machine. Nonetheless, because there are few ways to model or manipulate the human immune system other than by carrying out large-scale experiments on humans themselves, the amenability of this particular threat to improvement via the Design-Build-Test cycle is minimal, and predictability of results is likely to remain a significant barrier in the near term. Therefore, there is a medium-low level of concern with regard to this factor with the engineering of delivery systems amenable to delivery of immunomodulatory factors an area to monitor.

Requirements of Actors (Low Concern)

The expertise required to modulate human immunity with any degree of surety is likely quite high. In particular, choosing appropriate animal models for testing immunomodulatory interventions remains an art with only a few capable practitioners (Taneja and David, 2001; Benson et al., 2018). Moreover, several of the approaches considered would require an actor to not only successfully develop and deploy the immunomodulatory weapon itself but to successfully plan and execute a multipronged attack in which the immunomodulatory weapon is combined with another biological attack (such as deploying a pathogen after an initial attack causing immunodeficiency) or

specialized public health knowledge (such as vulnerabilities created by vaccination patterns, see Chapter 7, Health-Associated Data and Bioinformatics). Such approaches therefore increase the already advanced level of expertise required to effect an immunomodulatory attack, leading to an overall low level of concern for this factor. However, fast-advancing research in immunotherapies may reduce some of these barriers and expand the availability of the appropriate knowledge and skills in the coming years.

Potential for Mitigation (High Concern)

Modulation or evasion of the human immune system is already a hallmark of many pathogens, many of which are constantly developing novel means to avoid immune surveillance (e.g., seasonal adoption of new glycosylation sites by influenza) (Tate et al., 2014). There are also likely many unknown or undercharacterized pathogens that are currently biasing immune responsivity. These natural dynamics would make differentiating between natural and synthetic threats a considerable challenge. It may be particularly daunting to identify the hand of a designer versus the opportunism of nature in a given epitope in a pathogen variant that leads to autoimmunity. The lack of knowledge regarding the mechanisms for discriminating self versus non-self would also increase the challenges associated with recognizing an attack and deploying effective countermeasures. For these reasons, there is a relatively high level of concern with regard to this factor.

Whereas public health measures can potentially be useful in countering a threat involving immunomodulation, recognizing a problem and deploying the appropriate countermeasures would not necessarily be easy or quick; the slow response to the AIDS epidemic, albeit almost 40 years ago, is a potential cautionary tale in this regard. The current state of knowledge regarding immunity is such that it is likely far easier to craft an immunomodulatory weapon than an effective response to one. Even if good countermeasures could be crafted, their expense would likely be inordinate, especially for more general attacks on population immunity.

MODIFYING THE HUMAN GENOME

In addition to using synthetic genes to impact human physiology through pathogens or modifications to the microbiome, it may also be possible to insert engineered genes directly into the human genome via horizontal transfer, in other words, to use “genes as weapons.” Recent improvements in the ability to deliver genetic information via horizontal transfer, for example, through tools such as CRISPR/Cas9, potentially open the way for synthetic or cross-species transfer of genetic information into human hosts. In addition to protein-encoding genes, genes that encode RNA products such as short hairpin RNAs (shRNAs) or miRNAs could potentially be exploited as weapons in their own right. In combination with technologies for the modification of genes or their expression, deepening insights into systems biology could open new opportunities for causing diseases that are outside the rubric of the types of threats typically focused on in biodefense. Several ways in which synthetic biology approaches could be used to horizontally transfer genetic information to a human target to cause harm were considered:

- *Deletions or additions of genes.* If researchers can create mouse models of particular disease states based on the deletion or addition of particular genes, it follows that if the genomes of human beings could be similarly modified, such modifications could potentially cause a wide variety of noninfectious diseases. In particular, decades of research on genes associated with oncogenesis—oncogenes—have yielded many examples of gene changes that lead to cancer, including via infection by viruses and bacteria (Robinson and Dunning Hotopp, 2014; Cui et al., 2015; Sieber et al., 2016). Oncogenes could potentially be horizontally transferred to human cells via unnatural means. In this vein, CRISPR/Cas9 has been used to create point mutations, deletions, and complex chromosomal rearrangements in germline and somatic cells to develop mouse models for cancer (Mou et al., 2015).
- *Epigenetic modifications.* Just as programmed genetic modifications are possible, it may also prove possible to use horizontal transfer to alter the epigenetic state of an organism in a way that causes harm. Epigenetic modifications are clearly of immense importance in gene expression and are implicated in disease states and pathogenicity. For example, it is now proving possible to predict the course of oncogenesis based on

the epigenetic state of a tumor (Jones and Baylin, 2007). Sequence-specific epigenetic modifications can be carried out by small RNAs in other species, such as plants, but are not extensive in humans (He et al., 2011). However, the sequence-specific binding capabilities of Cas9 and other CRISPR elements may allow fusion proteins to carry out sequence-specific epigenetic modifications (Brocken et al., 2017). There are also chemicals that yield relatively nonspecific epigenetic changes (Bennett and Licht, 2018).

- *Small RNAs.* Small RNAs are another example of functional genetic information that could be horizontally transferred. Small RNAs, although not a genome modification per se, are important because they may prove capable of modifying gene expression and bringing about phenotypic change. The large number of small interfering RNA (siRNA), short hairpin RNA (shRNA), micro RNA (miRNA) (Zhang et al., 2007; Huang et al., 2008), and other small-RNA library studies in a variety of species and cells from different species, including human, provides a potential roadmap of what sequences may lead to what disease states or to modulation of defenses against disease. Similarly, there are already numerous viral and other vectors that can encode and express small RNAs. The fact that many viral pathogens already seem to encode small RNAs that aid in their pathogenicity further underlines this possibility. For example, the oncogenic gamma herpesviruses Epstein-Barr virus (EBV) and Kaposi's sarcoma-associated herpesvirus (KSHV) encode miRNAs that clearly act as mediators of immune suppression (Cullen, 2013). While most gene delivery mechanisms would likely be facilitated by CRISPR elements, direct delivery of small RNAs via liposomes or other vehicles has proven possible in many cell types (Barton and Medzhitov, 2002; Wang et al., 2010; Miele et al., 2012), and more recently the delivery of entire messenger RNAs (mRNAs) has proven useful for vaccination and cellular reprogramming (Steinle et al., 2017). Naked RNA is generally considered to be fragile due its susceptibility to ribonuclease in the cell, and its delivery is largely confined to laboratory settings, but there are approaches for stabilizing RNAs (e.g., using liposomes, nanoparticles, synthetic polymers, cyclodextrins, ribonucleoproteins, and viral capsids ["armored" RNAs]) in use for many applications. RNA can be expressed from genes delivered as simple expression vectors, as low-fitness-burden cargoes on viral pathogens, or via CRISPR element insertion. One reason that RNA delivery is potentially a viable biological threat is that even a small initial skew in gene expression (such as the changes in gene expression normally caused by miRNAs) could greatly alter the probability of an initial cellular alteration. Even small amounts of a targeted RNA would not modify the genome per se, but might allow or encourage cells to begin the process of self-transformation to tumors, as evidenced by the fact that a large number of pro-oncogenic miRNAs have already been discovered (O'Bryan et al., 2017). In addition to RNAs produced by viruses, bacteria produce numerous small regulatory RNAs; introduction of these into the endogenous microbiome could lead to dysbiosis. Larger mRNAs can also be delivered via liposomes and nanoparticles or by RNA replication strategies being developed for vaccine production (see Chapter 8, Rapid Development of Self-Amplifying mRNA Vaccines); these methods could potentially be used to express deleterious cargo such as toxins or oncogenes, similar to threats related to DNA vectors.
- *CRISPR/Cas9.* CRISPR elements can be harnessed for site-specific cleavage of genes, followed by homologous recombination via double-strand break repair or other mechanisms. This technology has revolutionized genome engineering. The fact that DNA recognition can be programmed by simple modification of an RNA element makes precision targeting of genome change much easier than previous technologies such as zinc finger endonucleases and TAL effector nuclease (TALEN)-mediated sequence-specific recognition of DNA. Another advantage of CRISPR technology is its broad host range; CRISPR elements are able to recognize and bind to DNA sequences in species other than those in which they originally evolved. Thus, the fact that gene editing technologies such as CRISPR make possible genomic changes in animal models that directly impact health and pathogenesis further implies that it may be possible to manipulate either germline or somatic cells to make such changes in humans. Significantly, the sequence specificity of CRISPR elements might also make possible ethnospecific targeting of gene-based weapons depending on the distributions of alleles (see also Chapter 7, Health-Associated Data and Bioinformatics). In terms of delivery, CRISPR elements could potentially be loaded onto a pathogen or delivered via the microbiome to modify human genomes in a way that would pose harm to individuals or populations.

- *Human gene drives.* Because of the ability of CRISPR elements to modify genomes, they can be repurposed as selfish genetic elements in their own right, wherein their introduction into a naïve genome leads to their site-specific establishment. In sexually reproducing organisms, an appropriately modified CRISPR element or other homing endonuclease gene, when used as a gene drive, can spread throughout a population. Gene drives are well known in nature, such as the *Drosophila* P element, which moves nonspecifically through naïve populations based on sexual (vertical) transfer. Gene drives have recently proven to be extremely useful for engineering mosquito populations for infertility (Hammond et al., 2016) and they have been proposed for the attenuation of fitness in other undesirable species, as well (for more detail, see National Academies of Sciences, Engineering, and Medicine, 2016). Concerns related to the use of gene drives in human populations were assessed separately from other potential approaches involving horizontal gene transfer because fundamental differences in the mechanisms involved in these different types of activity engender significantly different levels of concern. The assessment of concerns related to the use of human gene drives is summarized below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for modifying the human genome using human gene drives | Low | | | |

Assessment of Concerns About Gene Drives

For a gene drive to spread in a population, typically many cycles of reproduction are required so that genes can be vertically transferred from one generation to the next. Because humans have a relatively long generation span due to our age of reproductive maturity, a gene drive would take thousands of years to spread throughout a human population in this manner. In addition, some resistance mechanisms to gene drives are already becoming apparent as barriers to their use (Champer et al., 2017). In short, because of the fundamental and insurmountable constraint of human reproductive cycle length, the level of concern with regard to human gene drives is very low and other factors beyond usability of the technology were not analyzed.

The assessment of concerns related to modifications to the human genome through approaches *other than* through gene drives is summarized here and described in detail below.

| | Usability of the Technology | Usability as a Weapon | Requirements of Actors | Potential for Mitigation |
|---|-----------------------------|-----------------------|------------------------|--------------------------|
| Level of concern for modifying the human genome | Medium-low | Low | Medium-low | High |

Assessment of Concerns About Genome Modifications Other Than Gene Drives

Usability of the Technology (Medium-Low Concern)

Engineering genes to infiltrate an individual’s genome and cause harm is likely to be a technically challenging endeavor, leading to a medium-low level of concern with regard to this factor. Approaches focused on transient horizontal transfer of genes or small RNAs (e.g., via modified viral vectors) could be used, along with systems biology insights, to engineer changes in genes or gene expression to cause noninfectious disease, such as cancer

or neurological debilitation, or to degrade immunity. For example, the use of engineered pathogens to deliver small RNAs that cause healthy cells to initiate tumors may be feasible with current knowledge and technology. However, there would be significant challenges to determining the right targets or edits, packaging the genetic cargo into viral vectors, and delivering it to appropriate host cells.

CRISPR-based genome editing technologies are advancing rapidly and could be used to create genetic modifications propagated through engineered pathogenic vectors or horizontal transfer to human cells. However, it would likely be difficult to implement such genome modifications, in part because of the size of the protein-based machinery required for DNA recognition and cleavage, which would impose a hefty fitness cost on the (likely viral) pathogen unless it is linked with the viral life cycle in some way. In other words, viral pathogens have no need to cleave genomes, and this would likely limit the viability of viruses carrying genome-cleaving machinery. That said, new alternatives to the ubiquitous CRISPR/Cas9 system, such as the smaller Cpf1 (Zetsche et al., 2015), *Staphylococcus aureus* Cas (Ran et al., 2015), or newly discovered CasX and CasY (Burstein et al., 2017) could reduce this barrier.

If an actor sought to cause cancer in targeted individuals, it might only be necessary to modify a small number of cells to initiate oncogenesis and cause a self-sustaining and potentially metastatic cancer. Thus, the mechanisms for delivery could be relatively inefficient and might not require a replicating pathogen for initial distribution. A sufficient gene modification could be accomplished, for example, by introducing the ribonucleoproteins (RNPs) of CRISPR elements by themselves, rather than as genes, with an accompanying protein translocation domain to transit cellular membranes (Liu et al., 2015; Kouranova et al., 2016). This makes a CRISPR RNP potentially more akin to a toxin than to a traditional pathogenic biological threat. Similarly, DNA need not replicate to lead to expression in cells; there are many circular and linear plasmid vectors that can be transiently transfected into a host and thereby provide transient expression of even a large cargo (Nafissi and Slavcev, 2012). This route could be used to facilitate delivery of CRISPR/Cas9 and accompanying oncogenic guide RNAs to a host. In addition, a number of RNA-based mechanisms for gene delivery have come to the fore as a result of recent thrusts to create RNA-based vaccines (Kranz et al., 2016; Pardi et al., 2017). These methods lead to amplification of the originally introduced nucleic acid, but do not otherwise spread between individuals. Thus, they could be used to facilitate oncogenesis in a specifically targeted population.

Usability as a Weapon (Low Concern)

Even were it to become more technologically feasible to use genes to cause oncogenesis, neurodegenerative disease, immunological collapse, or other undesirable states, in the absence of a pathogen or greatly advanced unnatural horizontal transfer mechanism to promote the dispersal of a gene, the ability of an actor to deliver genes for these purposes is limited. Therefore, given this barrier, the concern level regarding usability as a weapon is relatively low. The mechanisms of dispersal (other than pathogens themselves) are likely to be low yield, the probability of inculcation of the disease state is likely to be low, and the onset of the disease state is likely not rapid. However, these limitations do not necessarily preclude an actor from pursuing such a weapon, especially since such a weapon could still significantly impact morale and readiness. In addition, many of these envisioned genetic weapons would become substantially more insidious if the skin rather than the bloodstream could be utilized as a route of entry, and improvements in dermal delivery could greatly change the landscape of threat. The use of siRNAs as a means of targeting tyrosine hydroxylase or tyrosinase and thereby treating hyperpigmented scars (Xiu-Hua et al., 2010) is instructive as to how this route may be actionable; it will be important to monitor future developments in this area.

Requirements of Actors (Medium-Low Concern)

Almost all of the technologies that might be instrumental in the use of genes as weapons are still in their translational infancy, practiced primarily in research laboratories and not in the clinic. Therefore, the concern level with regard to requirements of actors is medium-low. Achieving the types of potential bioweapons envisioned would likely require advanced research knowledge and experience, not just technical ability. Even advanced

companies that would be best suited for the development of dual-use technologies, such as siRNAs, have yet to fully develop delivery methods for desired biomedical applications. One possible exception is the development of bioweapons designed to cause cancer; possible approaches for such an attack can be inferred from knowledge of how chemicals in the environment have impacted cancer epidemiology and from laboratory data on how to induce cancers in animals. An additional caveat is that the rapid spread of technologies for genome engineering via CRISPR element toolsets could potentially decrease the barrier to entry for actors. For example, gene editing could be used to engineer a gene drive into an endemic insect or other pest population to assist delivery of a noxious or infectious agent. In this scenario, even a poorly functioning gene drive might not have to be successful for very long to achieve an effect.

Potential for Mitigation (High Concern)

Overall, the relative level of concern related to the potential for mitigation of gene-based weapons is high. Although some types of impacts would be readily recognized and attributed to a purposeful attack, it would be extremely difficult to trace some impacts—an epidemic of new cancers, for example—to a bioweapon. Such an attack may unfold very slowly, gradually skewing the health of a population. This would make mitigation very difficult, as presaged by experiences with identifying, tracing, and addressing cancer epicenters near toxic waste sites over the past several decades. The considerable challenge of mitigating an intentional cancer epidemic is a primary driver for the high level of concern relating to mitigation for this potential threat. However, once a threat is recognized, established mitigation methods such as quarantine and potential new ones such as therapeutic genome editing could be effective against some types of gene-based weapons.

Given that exome sequence data are being generated at an exponential rate, the introduction of CRISPR elements in humans or other higher organisms would likely be identified quickly and immediately recognized as cause for alarm. The presence of previously unknown oncogenes in viruses not normally known to harbor oncogenes would also be an immediate cause for alarm. However, the surreptitious spread of an oncogenic small-RNA sequence, especially if it is embedded within a protein-encoding gene, might be less noticeable and thus evade detection.

SUMMARY

- The alteration of humans through mechanisms that are different than conventional pathogens is an important potential concern area. The reduction or removal of key bottlenecks and barriers in the future could make some of the approaches discussed in this chapter more feasible.
- As understanding of microbiomes increases, the possibility of misuse also increases, and it may become feasible to use synthetic biology to engineer the microbiome to transfer toxic genes, debilitate human immunity, improve pathogen entry or spread, or create dysbioses.
- The threat posed by human immune modulation is limited by current knowledge, but knowledge is accumulating rapidly enough that it may well become more feasible to predictably modify the human immune system.
- Strategies to modify the human genome or alter gene expression in undesirable ways include gene editing, delivery of RNA molecules, and use of chemicals with epigenetic effects, although significant technical and delivery barriers remain that constrain feasibility.

While the traditional biodefense paradigm places agents such as pathogens or chemicals at the center of considerations of threat and vulnerability, this chapter attempts to reshape that paradigm by considering how interplay with and potential modifications of the human host might change the threat landscape. As understanding of the human microbiome, human immunity, and the human genome increases, the possibility of misuse also increases. In addition, advances in the understanding of individual genetic variability and in the ability to exploit individual

variation may make it more feasible to target host-modifying attacks to individuals or subpopulations (further discussed in Chapter 7, Health-Associated Data and Bioinformatics).

The current state of knowledge of the human microbiome is rapidly increasing, and it may be feasible to use synthetic biology to engineer the microbiome to transfer toxic genes, debilitate human immunity, improve pathogen entry or spread, or create dysbioses. However, with the exception of the *in situ* production of a hazardous compound (as detailed in Chapter 5, Making Biochemicals Via *In Situ* Synthesis), these potential threats are of lesser concern than more traditional pathogen- and chemical-centered attacks. Despite being an active area of research, the microbiome is still not fully understood, and creating a microbe that could colonize and persist within an established commensal community is a significant challenge. Furthermore, the judicious use of antibiotics could be an effective countermeasure to attacks propagated through the microbiome. Indeed, given the strong push to improve human health via microbiome research and engineering, there may be far more robust opportunities for microbiome-based countermeasures than threats.

The overall concern posed by human immune modulation is similar to the overall concern posed by microbiome engineering, and for similar reasons. On the one hand, current knowledge limitations likely preclude this potential vulnerability from being exploited in a significant way in the near future. On the other hand, knowledge is accumulating at such a rapid clip that it may well become more feasible to predictably modify the human immune system, and the expertise needed to do so is likely to become more widespread in the coming years. In addition, even unpredictable modifications can still cause harm. While it could have been predicted that IL-4 insertion into the mousepox genome would lead to the virus's ability to overcome vaccination (Müllbacher and Lobigs, 2001), it is still unknown whether the same type of modification in a human variant of a virus would have similar dire consequences. In contrast, the development of an anti-CD28 antibody was judged safe enough based on the rigorous review accorded clinical trials, yet proved to be life-threatening (Suntharalingam et al., 2006). Overall, the engineering of hyperimmunity and subsequent pathogenesis seems a greater threat than the engineering of reduced immunity or autoimmunity. The former is acute and fits more readily with individual pathogens and weaponization scenarios; the latter are chronic and with enough foresight can potentially be dealt with at a societal level via the usual public health measures for containing communicable diseases.

Building on that analysis, while the assessment focused on the human immune system, it is important to keep in mind that there are other potential systems that may also prove to be vulnerable to manipulation. For example, human neurobiology is immensely complex, and there are already a variety of genetic and chemical means to manipulate the overall mental health of individuals. That said, it is difficult to engineer such systems for a particular outcome with any surety. It will be important to continue to monitor advances related to understanding and modifying these complex systems in the coming years.

The concept of genes as weapons encompasses the development of synthetic genes that could change human physiology, either on their own or potentially delivered as an augment to a known pathogen. This concept also encompasses the possibility of delivering synthetic genes for small RNAs (or the synthetic small RNAs themselves) that could impact host physiology via interference mechanisms. Genes have a unique position in the biological threat pantheon, being somewhere between pieces of genomes, in which case they can be considered as just parts of pathogens, and being toxins, chemical compounds capable of harm without necessarily replicating. There are multiple difficulties that surround their delivery and a limited number of military scenarios in which an adversary would find it worthwhile to alter human physiology over time frames longer than a single battle or campaign. That said, some scenarios, such as the use of dermal transfection to create shRNAs or miRNAs that alter human physiology, or the use of gene drives to alter insect populations to deliver noxious compounds to humans, may present more attractive options from the perspective of an adversary.

In addition, threats related to horizontal gene transfer in synergy with the threats posed by pathogens may lead to new modes of attack. Just as clinical trials of immunotherapies are increasingly a roadmap for engineering cytokine storms, the increasing knowledge on gene deletions, gene additions, and small-RNA modifications of human cells may provide a roadmap for the induction of noninfectious disease states that could be abetted by pathogen engineering (and, conversely, that could abet the spread of the pathogens themselves, such as via immunodeficiency viruses).

Relevant developments to monitor for each of these capabilities are summarized in Table 6-1.