



**Office of Sponsored Programs**

4400 University Drive, MSN 4C6, Fairfax, Virginia 22030  
Phone: 703-993-2988; Fax: 703-993-2296)

May 3, 2016

Open Philanthropy  
182 Howard Street #208  
San Francisco, CA 94105

Dear Contracting Officer,

Enclosed please find a submission on behalf of Dr. Robin Hanson, Associate Professor, Economics, Center for Study of Public Choice, George Mason University. The project is entitled "AI As Software Broad Analysis."

Mason believes the project proposed herein is fundamental research as defined in National Security Decision Directive 189, and to our knowledge, does not require that we obtain an export license under EAR or ITAR. If Open Philanthropy believes that our performance or deliverables under this project are subject to export control regulations, and therefore not eligible for the fundamental research exclusion, we request written confirmation to that effect.

This proposal has been reviewed and administratively approved by the appropriate university officials. For questions regarding the technical aspects of this proposal, please feel free to contact Dr. Robin Hanson at (703) 993-2326. Any questions regarding budget or university policies and procedures should be directed to me at (703) 993-4806 or [ospaor@gmu.edu](mailto:ospaor@gmu.edu).

Sincerely,

A handwritten signature in cursive script that reads "Donna Senator".

Donna Senator  
Associate Director, Proposal & Award Management  
Office of Sponsored Programs

Enclosures

cc: Dr. Robin Hanson, Principal Investigator

# George Mason University - Proposal Budget Summary

Coeus Number: 00006844

Proposal Number: 116989

Proposal Title: AI As Software Broad Analysis

Investigator: Hanson, Robin D

Sponsor: Open Philanthropy

		Effort%	Period 1 05/25/2016 to 05/24/2017	Period 2 05/25/2017 to 05/24/2018	Period 3 05/25/2018 to 05/24/2019	Total
<b>Senior Personnel</b>						
Hanson, Robin D	Faculty Salary: Academic or Calendar Year	25	\$20,813.41	\$21,437.81	\$22,080.94	\$64,332.16
Hanson, Robin D	Faculty Summer Salary	50	\$13,471.46	\$13,875.61	\$14,291.87	\$41,638.94
<b>Other Personnel</b>						
	Wages or Admin Staff Salaries		\$30,078.00	\$30,980.34	\$31,909.75	\$92,968.09
<b>Fringe Benefits</b>						
			\$10,068.36	\$10,370.40	\$10,681.51	\$31,120.27
			<b>Subtotal Personnel + Fringe:</b>	<b>\$74,431.23</b>	<b>\$78,964.07</b>	<b>\$230,059.46</b>
<b>Travel - Domestic</b>						
	Domestic Travel		\$1,000.00	\$1,030.00	\$1,060.90	\$3,090.90
<b>Travel - Foreign</b>						
	Foreign Travel		\$1,400.00	\$1,442.00	\$1,485.26	\$4,327.26
			<b>Subtotal Travel:</b>	<b>\$2,400.00</b>	<b>\$2,546.16</b>	<b>\$7,418.16</b>
<b>Other Operating Expenses</b>						
	Other Direct Expenditures		\$3,000.00	\$0.00	\$0.00	\$3,000.00
			<b>Subtotal Other Operating Expenses:</b>	<b>\$3,000.00</b>	<b>\$0.00</b>	<b>\$3,000.00</b>
			<b>Total Direct Costs:</b>	<b>\$79,831.23</b>	<b>\$79,136.16</b>	<b>\$81,510.23</b>
			<b>F&amp;A (Indirect) Costs:</b>	<b>\$7,983.13</b>	<b>\$7,913.61</b>	<b>\$8,151.03</b>
			<b>Total Costs:</b>	<b>\$87,814.36</b>	<b>\$87,049.77</b>	<b>\$89,661.26</b>

**George Mason University**  
**Robin Hanson / AI As Software Broad Analysis**  
**Proposal 116989**

**BUDGET JUSTIFICATION**

**PERSONNEL**

Faculty

Professor Robin Hanson will serve as the PI of this proposal at George Mason University. Dr. Hanson will dedicate 2.25 academic months and 1.5 summer months to administer the project.

A TBD non Student Wage will be hired. Total requested in year one is \$30,078.00.

**MERIT INCREASES**

Mason provides annual merit increases to Faculty and Staff. An escalation factor of 3% has been included for all personnel each year.

**FRINGE BENEFITS**

George Mason University's negotiated fringe benefit rates for Fiscal Year 2016 are applied as follows:

Faculty (Admin, Teaching, & Post-Docs)	33.1%
Classified Staff	42.9%
FICA Only (summer, adjunct, non-student wages)	7.3%
Student wage	6.9%

\* Student wage employees taking a full-time class load are generally exempt from FICA. Student wage employees taking less than a full-time class load are not eligible for full-time student FICA exemption and will be charged the student fringe benefit rate of 6.9%.

Salaries, wages and fringe benefits are estimates only and will be paid and billed in accordance with University policy.

**TRAVEL**

\$1000 per year is estimated for domestic travel to technical conferences to present the results of our research. This represents 1 trip per year. An additional \$1400 is estimated for foreign travel to 1 technical conference per year. Travel lines are inflated by 3% in years two and three. This budget line is based on the following assumptions:

Round trip airfare in US:	1 trip x 1 staff x \$610 per flight.....	\$610
Round trip airfare international:	1 trip x 1 staff x \$1010 per flight.....	\$1010
Lodging for both trips:	1 trip x 2 night stay x 1 staff x \$120.....	\$240
Meals and other per day for both trips:	1 trip x 2 days x 1 staff x \$75.....	\$150

**SUPPLIES**

The proposal's budget includes \$3000 in the first year, \$0 in later years, for a new computer and supporting software.

**FACILITIES AND ADMINISTRATIVE COSTS (F&A)**

This sponsor requires that indirect costs be at a rate no higher than TDC 10%. Total indirect costs for this project are \$24,048.

# Detailed Description of “AI as Software Broad Analysis”

Robin Hanson of George Mason University, Proposal to Future of Life Institute, May 17, 2015.

## Introduction

In AI safety, as elsewhere, solutions should depend on context. Yes, we prefer robust solutions, but robustness can be expensive. Because of this, learning more about the contexts of particular problems can often help greatly to more efficiently target solutions. Thus those who seek to mitigate AI risk should want to learn more about how AI will actually develop and then impact society.

We can divide up this learning task in (at least) four different ways. First, we can distinguish learning about paths that lead to AI from learning about consequences after AI has arrived. Second, we can distinguish “unipolar” scenarios, wherein a single actor or tight coalition dominates outcomes, from “multipolar” scenarios, wherein many independent actors are influential. Third, we can distinguish different types of AI that might first achieve human-level abilities. For example, such AIs might be based on emulations of human brains, on a relatively steady accumulation of software tools, or on a more dramatic breakthrough in mind design architecture. Fourth, we can do positive analysis, predicting likely outcomes, or we can do normative analysis, designing and recommending some choices relative to others.

Robin Hanson proposes to take three years to conduct a **broad positive analysis** of the **multipolar scenario** wherein AI results from relatively steady **accumulation of software tools**. That is, he proposes to assume that human level AI will result mainly from the continued accumulation of software tools and packages, with distributions of cost and value correlations similar to those seen so far in software practice, in an environment where no one actor dominates the process of creating or fielding such software. He will attempt a mostly positive analysis of the social consequences of these assumptions, both during and after a transition to a world dominated by AI. While this is hardly the universe of all desired analyses, it does seem to cover a non-trivial fraction of interesting cases.

## Qualification

Prof. Hanson seems unusually well qualified to conduct such an analysis. He spent nine years writing artificial intelligence software for a living, and for sixteen years has been a professor in economics, a discipline that focuses mainly on positive analysis of multi-polar social scenarios. Finally Prof. Hanson has unusually broad expertise, having published in computer science, physics, philosophy, economics, and political science. Google Scholar says he has 2720 citations, roughly equally divided between economics and computer science.

More importantly, Prof. Hanson has demonstrated relevant abilities in having just successfully produced a closely related analysis. His main project for the last few years has

been a positive analysis of the post-transition social consequences of a multi-polar brain-emulation-based AI scenario. This analysis has been mostly opportunistic and qualitative. Drawing on many academic disciplines, he applied as many simple standard analysis tools as he could make relevant. He has arguably achieved dramatically unprecedented breadth and detail, elaborated in a book *The Age Of Em: Envisioning Brain Emulation Societies* to be published by Oxford University Press by spring 2016. A 260-page current book draft is available for review. Here is the outline:

## **I. Basics**

1. *Start*: Contents, Preface, Introduction, Summary
2. *Modes*: Precedents, Factors, Dreamtime, Limits
3. *Mechanics*: Emulations, Opacity, Hardware, Security

## **II. Physics**

4. *Scales*: Time, Space, Reversing
5. *Infrastructure*: Climate, Cooling, Buildings
6. *Existence*: Virtuality, Views, Fakery, Copying, Darkness
7. *Farewells*: Fragility, Retirement, Death

## **III. Economics**

8. *Labor*: Wages, Selection, Enough
9. *Efficiency*: Competition, Eliteness, Spurs, Power
10. *Business*: Institutions, Growth, Finance, Manufacturing
11. *Lifecycle*: Careers, Age, Preparation, Training

## **IV. Organization**

12. *Clumping*: Cities, Speeds, Transport
13. *Extremes*: Software, Inequality, War
14. *Groups*: Clans, Nepotism, Firms, Teams
15. *Conflict*: Governance, Law, Innovation

## **V. Sociology**

16. *Connection*: Mating, Signaling, Identity, Ritual
17. *Collaboration*: Conversation, Synchronization, Coalitions
18. *Society*: Profanity, Divisions, Culture, Stories
19. *Minds*: Humans, Unhumans, Intelligence, Psychology

## **VI. Implications**

20. *Variations*: Trends, Alternatives, Transition, Aliens
21. *Choices*: Evaluation, Policy, Charity, Success
22. *Finale*: Critics, Conclusion, References, Thanks
23. *Appendix*: Motivation, Method, Biases

Brain emulations have been a staple of science fiction and futurism for decades, and many have eagerly speculated on the timing of this technology, on whether such emulations would be conscious, and whether they would be “the same person” as the original scanned and emulated human. However, prior to this new book, there has been relatively little careful expert analysis of the social implications of brain emulations.

## Leverage

This prior work demonstrates both the feasibility of Prof. Hanson's approach, and his expertise in implementing such an approach. If funded by the Future of Life Institute for three years, Hanson proposes to apply this method and expertise to a second and distinct AI scenario: human level AI based on the steady accumulation of software tools. In this scenario, the cost and productivity of AI software are assumed to be usefully described as typical of some already observed types of software.

No individual or organization intentionally funded Prof. Hanson's prior effort to analyze the social consequences of emulation-based AI. This effort required few resources beyond time, and as a tenured professor Hanson has been free to allocate most of his time as he chooses. Since this effort has not yet inspired interest among patrons of research, it seems unlikely that any other such patrons will fund related research anytime soon.

Furthermore, Prof. Hanson has many other attractive research agendas competing for his attention. While he could proceed with this research agenda even without funding, more likely he would focus on one of his other research areas. And at the moment no one else known to him seems inclined to pursue this proposed research agenda. Thus if the Future of Life Institute does not fund this new line of research on AI as like familiar software, it most likely will not be funded, and will not be pursued, by anyone, at least for a while.

## Strategy

Prof. Hanson proposes to follow this basic research strategy: *collect, apply, and refine*. Prof. Hanson proposes that during his first year, he will **collect** the best available models describing relevant software engineering cost and productivity. Such models explain how the cost and value of software depends on parameters like task range, number of previous efforts, and task environment variability. Costs include those to write, test, and augment code, and to adapt it to new environments. Related models describe the productivity variance of software engineers and their tools across tasks and cases. When software specific models are not available, he will try to substitute economic models from related industries.

By the second year of this project, Prof. Hanson proposes to choose ways to describe near-human-level AI systems in terms of the parameters of these collected models. This will allow him to **apply** these models to the case of the value and of costs of AI software systems, especially regarding the task of software engineering in similar systems. Such models, combined with our standard understandings of larger technical and social systems, should allow Prof. Hanson to forecast many aspects of development and consequences of human level AI in this scenario. Such forecasts may plausibly be analogous in number and detail to those already found in the brain emulation scenario.

Prof. Hanson has many years of experience building and applying formal mathematical models of technical and social systems. However in his work on the brain emulation scenario, he found that a mostly qualitative breadth-first exploration of possible consequences seemed to be the most cost-effective way to produce relevant insight. Qualitative analysis can quickly grab low hanging fruit, while formal models are typically

slow and narrow by comparison. Prof. Hanson expects the same to apply when studying AI as software. While there should eventually be an important place for formal models in this area, the first priority is to try to ensure that important basic considerations have not been neglected. Yes, some important considerations may require formal models to make them clear, but the low-hanging fruit for analysis are the considerations that are already clear without such models.

By the third year of this project, Prof. Hanson proposes to begin to **refine** his applications of collected models into forms suitable for wider communication. He will give publish and give talks, giving both detailed conclusions to technical audiences, and more accessible summaries to broader audiences. He may even write a book similar to his book on emulations. Prof. Hanson has demonstrated all of these communication abilities.

In a brain emulation based AI scenario, the transition to a world dominated by emulations seems especially jerky and disruptive, since partial or near emulations seem to have far less economic value. Since this makes it especially hard to forecast such a transition, in his prior analysis Prof. Hanson focused mostly on forecasting post-transition outcomes. However, relative to emulations AI as software below the human level would have far more economic value. This should make it easier to forecast the transition to a world dominated by AI as software. Because of this, Prof. Hanson may focus more in this work on pre- and during- transition outcomes, in addition to post-transition outcomes.

The final refined analysis of AI as software will be compared to the analysis of the AI via brain emulations scenario, and to other available analyses. Promising areas to continue in related lines of research may also be identified.

## **Conclusion**

Prof. Robin Hanson proposes to take three years to conduct a broad positive analysis of a multipolar scenario wherein AI results from relatively steady accumulation of software tools. This effort would build on his recent successful effort to conduct a broad positive analysis of a multipolar scenario wherein AI results from brain emulations. If the Future of Life Institute does not fund this analysis, such analysis will most likely not be done by anyone for a while.

# Robin D. Hanson

Associate Professor of [Economics](#), [George Mason University](#)  
Fairfax, VA 22030 [rhanson@gmu.edu](mailto:rhanson@gmu.edu) 703-993-2326  
Research Associate, Future of Humanity Institute at Oxford University  
Chief Scientist, Consensus Point

## **Five Representative Publications**

[Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation](#), *Journal of Prediction Markets* 1(1):3-15, February, 2007. (211 citations)

[Information Aggregation and Manipulation in an Experimental Market](#). with [Ryan Oprea](#), David Porter, *Journal of Economic Behavior and Organization* 60(4):449-459, August 2006. (153 citations)

[Combinatorial Information Market Design](#). *Information Systems Frontiers* 5(1):105-119, January 2003. (293 citations)

[Super-Resolved Surface Reconstruction From Multiple Images](#). with [P. Cheeseman](#), [B. Kanefsky](#), [R. Kraft](#), [J. Stutz](#), *Maximum Entropy and Bayesian Methods* 293-308, ed. G.R. Heidbreder, [Kluwer](#), 1996. (250 citations)

[Bayesian Classification with Correlation and Inheritance](#). with [J. Stutz](#), [P. Cheeseman](#), *Proceedings of the 12th International Joint Conference on Artificial Intelligence* 2:692-698. Morgan Kaufmann, 1991. (218 citations)

## **Five Relevant Publications**

[The Hanson-Yudkowsky AI-Foom Debate](#), book with Eliezer Yudkowsky, Berkeley, CA: Machine Intelligence Research Institute. 2013.

[Economics of the Singularity](#), *IEEE Spectrum*, 37-42, June 2008. Reprinted as: A New World Order, *Cosmos* 26:47-54, April/May 2009. (41 citations)

[Economic Growth Given Machine Intelligence](#), Working Paper, 2001. (19 citations)

[Burning the Cosmic Commons: Evolutionary Strategies of Interstellar Colonization](#), Working Paper, 1998. (22 citations)

[If Uploads Come First](#). 6(2):10-15 1994. (43 citations)

## **Education**

Seminars on Social Science & Health, American Medical Care System, 1997, U.C. Berkeley.

PhD 1998, [California Institute of Technology](#),

Advisors: [J. Ledyard](#), [R. McKelvey](#), [T. Palfrey](#), [S. Wilkie](#).

Dissertation: [Four Puzzles in Information and Politics: Product Bans, Informed Voters, Social Insurance, & Persistent Disagreement](#).

MS, MA ([physics](#), [philosophy of science](#)) 1984, [University of Chicago](#).

BS ([physics](#)) 1981, [University of California at Irvine](#).



## Experience

[George Mason University](#), Associate Professor of [Economics](#), 1999-present.

Taught undergraduate: Microeconomics, Health Econ., Law & Econ., Enviro. Econ., Urban Econ.

Taught graduate: Microeconomics, Industrial Organization

[University of California - Berkeley](#), [RWJF Scholar in Health Policy Research](#), 1997-1999.

[California Institute of Technology](#),

T.A. to [C. Plott](#), I. Lee, economics principles, to [R. Kiewiet](#), political science principles, 1995-1996.

R.A. to [J. Ledyard](#) & [D. Porter](#), institution design and experiments for FCC & NASA, 1993-1995.

[NASA Ames Research Center](#), research in [Bayesian statistics](#), 1989-1993.

[Xanadu Inc.](#), consultant on hypertext publishing design, 1988-1991.

[Lockheed Artificial Intelligence Center](#), research in machine learning, 1984-1989.

## Honors

Associate Editor, *Journal of Economic Behavior and Organization*, 2009+.

Associate Editor, *Journal of Prediction Market*, 2007+.

[Alfred P. Sloan](#) Dissertation Fellowship, 1996.

[Global IdeaBank Web Social Innovations Award, 1996.](#)

[Prix Ars Electronica Golden Nica, World Wide Web, 1995.](#)

[Institute for Humane Studies Fellowship](#), 1993.

NASA Space Act Award, ARC-12799-1, 1992.

Honorable Mention, NSF Fellowship, 1982.

## Peer Reviewed Publications

### Social Science Journal Articles

[Are Disagreements Honest?](#), with Tyler Cowen, *Journal of Economic Methodology*, to appear.

[Shall We Vote on Values, But Bet on Beliefs?](#), *Journal of Political Philosophy*, 21(2):151-178, June, 2013.

[Meet The New Conflict, Same As The Old Conflict](#), *Journal of Consciousness Studies* 19(1-2):119-125, 2012.

[Gaming Prediction Markets: Equilibrium Strategies with a Market Maker](#), with Yiling Chen, Stan Dimitrov, Rahul Sami, Daniel Reeves, David Pennock, Lance Fortnow, Rica Gonen, *Algorithmica*, 58(4):930-969, 2010.

[On Market Maker Functions](#), *Journal of Prediction Markets* 3(1):61-63, April 2009.

[A Manipulator Can Aid Prediction Market Accuracy](#), with [Ryan Oprea](#), *Economica*, 76(302):304-314,

April, 2009.

[An Experimental Test of Combinatorial Information Markets](#), with John Ledyard, Takashi Ishikida, *Journal of Economic Behavior and Organization* 69:182-189, 2009.

[Insider Trading and Prediction Markets](#), *Journal of Law, Economics, and Policy* 4(2):449-463, Spring 2008.

[The Promise of Prediction Markets](#), with Kenneth J. Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz, *Science* 320(5878):877-8878, May 16, 2008.

[Making Sense of Medical Paternalism](#), *Medical Hypotheses* 70(5):910-913, 2008.

[Showing That You Care: The Evolution of Health Altruism](#), *Medical Hypotheses* 70(4):724-742, 2008.

[The Policy Analysis Market: A Thwarted Experiment in the Use of Prediction Markets for Public Policy](#), *Innovations* 2(3):73-88, Summer 2007.

[The Hanson-Hughes Debate on "The Crack of a Future Dawn"](#), with James Hughes, *Journal of Evolution and Technology* 16(1):99-126, June 2007.

[Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation](#), *Journal of Prediction Markets* 1(1):3-15, February, 2007.

[Uncommon Priors Require Origin Disputes](#). *Theory and Decision* 61(4):318-328, December 2006.

[Information Aggregation and Manipulation in an Experimental Market](#). with [Ryan Oprea](#), David Porter, *Journal of Economic Behavior and Organization* 60(4):449-459, August 2006.

[Designing Real Terrorism Futures](#), *Public Choice* 128(1-2):257-274, July 2006. Also to appear in ed. Charles Rowley, *The Political Economy of Terrorism*, 2007, and ed. William Hancock, *Business Continuity and Homeland Security*, Edward Elgar, 2007.

[Adverse Selection in Group Insurance: The Virtues of Failing to Represent Voters](#). *Economics of Governance* 6(2):139-157, July 2005. Version in Ph.D. thesis, 1997.

[Warning Labels as Cheap Talk: Why Regulators Ban Drugs](#). *Journal of Public Economics* 87(9-10):2013-2029, September 2003. Version in Ph.D. thesis, 1997.

[For Bayesian Wannabes. Are Disagreements Not About Information?](#) *Theory and Decision* 54(2):105-123, March 2003. Version in Ph.D. thesis, 1997.

[Combinatorial Information Market Design](#). *Information Systems Frontiers* 5(1):105-119, January 2003.

[Disagreement Is Unpredictable](#). *Economics Letters* 77(3):365-369, November 2002.

[Why Health Is Not Special: Errors in Evolved Bioethics Intuitions](#). *Social Philosophy & Policy* 19(2):153-179, Summer 2002. Reprinted in *Bioethics*, ed. E.F. Paul, F.D. Miller, & J. Paul, Cambridge University Press, 153-179, 2002.

[How To Live In A Simulation](#). *Journal of Evolution and Technology* 7, September 2001. Reprinted in *Folha de S.Paulo* (a Brazilian newspaper) October 21, 2001. Covered in [10 press articles](#).

[Decision Markets](#). *IEEE Intelligent Systems* 14(3):16-19, May/June, 1999. Reprinted in *Entrepreneurial*

[Economics](#), ed. Alexander Tabarrok, Oxford University Press, 79-85, 2002.

[Consensus By Identifying Extremists](#). *Theory and Decision* 44(3):293-301, 1998.

[Is A Singularity Just Around The Corner?](#) *Journal of Evolution and Technology* 2, June 1998.

[Correction to McKelvey and Page, "Public and Private Information: An Experimental Study of Information Pooling"](#). *Econometrica* 64(5):1223-1224, 1996.

[Could Gambling Save Science? Encouraging an Honest Consensus](#). *Social Epistemology* 9(1):3-33, 1995. First appeared in *Gambling and Commercial Gaming: Essays in Business, Economics, Philosophy, and Science* 399-440, ed. W. Eadington & J. Cornelius, 1992. (Is proceedings of August 1990 conference.) Summary published, Russ Ray, Idea Futures A Free-Market Approach to Academic Research. *Futures Research Quarterly* 12(2):81-91. Summer, 1996.

[Comparing Peer Review to Information Prizes](#). *Social Epistemology* 9(1):49-55, 1995. See also [Reply to Comments](#). 9(1):45-48 1995.

[Buy Health, Not Health Care](#). *Cato Journal* 14(1):135-141, Summer 1994.

[Can Wiretaps Remain Cost-Effective?](#) *Communications of the Association of Computing Machinery* 37(12):13-15, December 1994. First appeared in *The Third CPSR Cryptography and Privacy Conference Sourcebook* 357-363, ed. David Banisar and Marc Rotenberg, June, 1993. Reprinted in [The Electronic Privacy Papers: Documents on the Battle for Privacy in the Age of Surveillance](#) 19-25, ed. B. Schneier & D. Banisar, John Wiley and Sons, 1997.

Comment on the scientific status of econometrics. *Social Epistemology* 7(3):255-256, 1993.

[Idea Futures: Encouraging an Honest Consensus](#), *Extropy* 3(2):7-17, Winter , 1992; See also [If Uploads Come First](#). 6(2):10-15 1994. [Nanarchy](#). 6(1):32-36, 1994. [Wormhole Warfare](#). 6(1):38-39, 1994. [Lilliputian Uploads](#). 7(1):30-31, 1995. [Great Idea Seeks Champion](#). 16:8, 1995. [A Critical Discussion of Vinge's Singularity Concept](#). Editor & Discussant, October 1998. Idea Futures and A Critical Discussion reprinted in *The Transhumanist Reader*, Wiley-Blackwell, ed. Max More and Natasha Vita-More, 2013, pp.243-257,395-418.

[Market-Based Foresight](#). *Foresight Update* 10:1,3-4, October 30, 1990. See also [Reply](#). 11:11, March 15, 1991. [Has Penrose Disproved A.I.? 12:4-5](#), August 1, 1991.

#### **Other Journal Articles**

[Drift-Diffusion in Mangled Worlds Quantum Mechanics](#), *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 462(2069):1619-1627, May 8, 2006.

[When Worlds Collide: Quantum Probability From Observer Selection?](#) *Foundations of Physics* 33(7):1129-1150, July 2003.

[Bayesian classification scheme](#). with John Stutz, Peter Cheeseman, Will Taylor, and Matthew Self, *Applied Optics* 31(27):5763 (in NASA Patter), September 20, 1992.

[Toward Hypertext Publishing, Issues and Choices in Database Design](#). *ACM SIGIR Forum* 22(1-2):9-26, Winter 1988.

#### **Books**

[The Hanson-Yudkowsky AI-Foom Debate](#), with Eliezer Yudkowsky, Berkeley, CA: Machine Intelligence Research Institute. 2013.

## Book Chapters

[A Tale Of Two Transitions](#), in *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*, ed. Ben Goertzel, Ted Goertzel, 2015.

[What Will It Be Like To Be An Emulation?](#) pp.298-309 in ed. Russell Blackford, Damien Broderick, *Intelligence Unbound: The Future of Uploaded and Machine Minds*, Wiley, August 18, 2014.

[Comments](#) on "Intelligence Explosion: Evidence and Import," pp.41-42, and [Comments](#) on "Some Economic Incentives Facing a Business that Might Bring About a Technological Singularity," p.159. in *Singularity Hypotheses: A Scientific and Philosophical Assessment*, ed. Amnon H. Eden, James H. Moor, Johnny H. Søraker, Eric Steinhart, Springer, 2013.

[Enhancing Our Truth Orientation](#). *Human Enhancement*, ed. Julian Savulescu and Nick Bostrom, Oxford University Press, March 2009.

[Catastrophe, Social Collapse, and Human Extinction](#), *Global Catastrophic Risks*, pp. 363-377, ed. Martin Rees, Nick Bostrom, and Milan Cirkovic, Oxford University Press, July 17, 2008.

[The Rapacious Hardscrapple Frontier](#), *Year Million: Science at the Far Edge of Knowledge*, pp. 168-192, ed. Damien Broderick, Atlas Books, May 19, 2008.

[Decision Markets for Policy Advice](#). *Promoting the General Welfare: American Democracy and the Political Economy of Government Performance*, 151-173, ed. Eric Patashnik and Alan Gerber, Brookings Institution Press, November 2006.

[Foul Play in Information Markets](#). *Information Markets: A New Way of Making Decisions in the Public and Private Sectors* 126-141, ed. Bob Hahn and Paul Tetlock, AEI Spring Press, May 2006.

[Fear of Death and Muddled Thinking -- It Is So Much Worse Than You Think](#). *Death And Anti-Death, Volume 3: Fifty Years After Einstein, One Hundred Fifty Years After Kierkegaard*, ed. Charles Tandy, Ria University Press, December 2005.

[Game Theory in Public Choice](#). *The Encyclopedia of Public Choice* II:258-261, eds. Charles Rowley and Friedrich Schneider, Kluwer Academic Publishers, 2004.

[Was Cypher Right? \(Part I\): Why We Stay In Our Matrix](#). *Taking the Red Pill: Science, Philosophy and Religion in the Matrix*, pp.23-32, ed. Glenn Yeffeth, BenBella Books, 2003. (Reviews: [1](#), [2](#), [3](#).)

## Published Conference Proceedings

Wei Sun, Hanson, R., Laskey, K.B. and Twardy, C. Trade-Based Asset Models for Combinatorial Prediction Markets. *Proceedings of the Eleventh UAI Bayesian Modeling Applications Workshop*. Quebec, Canada, July 2014.

[Combinatorial Prediction Markets: An Experimental Study](#), with Walter Powell, Robin Hanson, Kathryn Laskey and Charles Twardy, *Proceedings of the Seventh International Conference on Scalable Uncertainty Management*, Washington DC, September 16-18, 2013.

[Learning Parameters by Prediction Markets and Kelly Rule for Graphical Models](#), with Wei Sun, Robin Hanson, Kathryn B. Laskey, Charles Twardy, In *Proceedings of the First Big Data meet Complex Models Applications Workshop*, held at the Uncertainty in Artificial Intelligence Conference, Bellevue, WA, July 15, 2013.

[Probability and Asset Updating using Bayesian Networks for Combinatorial Prediction Markets](#), with Wei Sun, Kathryn Laskey, Charles Twardy, pp.815-824, ed. Nando de Freitas and Kevin Murphy,

*Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, Catalina Island, August 15-17, 2012.

[Economics of Brain Emulations](#), *Unnatural Selection - The Challenges of Engineering Tomorrow's People*, pp.150-158, ed. Peter Healey and Steve Rayner, EarthScan, London, December, 2008.

[Accept Interdependence](#), *Ci'num - The Digital Civilizations Forum 2006*, p.36, Aquitaine Europe Communications. 2006.

[Five Nanotech Social Scenarios](#). *Nanotechnology: Societal Implications II, Individual Perspectives*, ed. Mihail C. Roco and William S. Bainbridge, Springer, 109-113, November 2006.

[The policy analysis market: an electronic commerce application of a combinatorial information market](#). with Charles Polk, John Ledyard, Takashi Ishikida. *Proceedings of the ACM Conference on Electronic Commerce 2003*: 272-273. See also [poster](#).

[The Story of Idea Futures](#). with [Mark James](#), [Sean Morgan](#), *Prix Ars Electronica 95*, *International Compendium of the Computer Arts*, ed. H. Leopoldseder, C. Schopf, 54-59, 1995.

[Super-Resolved Surface Reconstruction From Multiple Images](#). with [P. Cheeseman](#), [B. Kanefsky](#), [R. Kraft](#), [J. Stutz](#), *Maximum Entropy and Bayesian Methods* 293-308, ed. G.R. Heidbreder, [Kluwer](#), 1996. Version of NASA Ames Tech Report FIA-93-02, February, 1993.

[Reversible Agents: Need Robots Waste Bits to See, Talk, and Achieve?](#) *Proceedings of Workshop on Physics and Computation: PhysComp '92*, 284-288. IEEE Computer Society Press, 1992.

[Bayesian Classification with Correlation and Inheritance](#). with [J. Stutz](#), [P. Cheeseman](#), *Proceedings of the 12th International Joint Conference on Artificial Intelligence* 2:692-698. Morgan Kaufmann, 1991. Version of NASA Ames Tech Report FIA-90-12-7-01, May, 1991.

[Research Patronage and Distributing A.I. Agents: Similar Problems with Similar Solutions?](#) *AAAI Fall Symposium on Knowledge and Action at Social and Organizational Levels* 68-69, 1991.

[Even Adversarial Agents Should Appear to Agree](#), *Proceedings IJCAI-91 Workshop On Reasoning In Adversarial Domains*, 1991.

## **Working Papers**

### **Requested Revisions**

[Eliciting Objective Probabilities via Lottery Insurance Games](#), revising for *Theory and Decision*.

[Long-Term Growth As A Sequence of Exponential Modes](#), revising for *Journal of Economic Behavior & Organization*.

[Patterns of Patronage -- Why Grants Won Over Prizes in Science](#), revising for *Economic Inquiry*.

[Burning the Cosmic Commons: Evolutionary Strategies of Interstellar Colonization](#), revising for *Icarus*.

[Must Early Life Be Easy? The Rhythm of Major Evolutionary Transitions](#), revising for *Origins of Life*.

[Economic Growth Given Machine Intelligence](#), revising for *Journal of Artificial Intelligence Research*.

### **Under Review**

[Price Elasticity of Demand in Employer-Provided Self-Insured Health Plans](#), with Iwona Kicingier

[The Determinants of the Quantity of Health Insurance: Evidence from Self-Insured and Not Self-Insured Employer-Based Health Plans](#), with Iwona Kicingier

[Can Manipulators Mislead Prediction Market Observers?](#), with Ryan Oprea, David Porter, Chris Hibbert, and Dorina Tila.

[Choosing Between Health and Wealth - A Survey](#), with Robert F. Graboyes and James W. Monks.

[The Informed Press Favored the Policy Analysis Market.](#)

[He Who Pays The Piper Must Know The Tune.](#)

[Location Discrimination in Circular City, Torus Town, and Beyond.](#)

[On Voter Incentives To Become Informed.](#) Version in Ph.D. thesis, 1997. Version as Caltech Social Science Working Paper 968, September 1996.

### **In Preparation**

[Forager Facts](#), with David Youngberg.

[When Do Extraordinary Claims Give Extraordinary Evidence?](#)

[Causes of Confidence in Conflict.](#)

[Book Orders for Market Scoring Rules.](#)

Priors Over Indexicals.

An Experimental Test of Agreeing to Disagree, with William Nelson.

[Is Fairness About Clear Fitness Signals?](#)

[World Peace, Thanks To Old Men?](#)

[The Great Filter - Are We Almost Past It?](#)

[Rational Bar Bets](#)

[Shared Secrets Come Cheap](#)

Should Privately Enforced Laws Fix Punishment or Restitution?

[Composing Bayesians: A Short Review](#)

### **Other Publications**

[Should Earth Shut the Hell Up?](#), [Pascal's Alien Wager](#), [Slow Growth Is Plenty Fast](#), [Adapted Aliens](#), [Selection Is Coming](#), *Cato Unbound*, December 2014.

Regulating Infinity, *Global Government Venturing*, pp.30-31, September 2014.

[When the Economy Transcends Humanity](#), *The Futurist*, January-February, 2014.

[Who Cares About Forecast Accuracy?](#), *Cato Unbound*, July 2011.

[Open-Source Government: How can we fix our political system?](#), *BBC Focus* 206:21, August 2009.

[LHC forecasts: better than horoscopes?](#), *Symmetry* 5(5):32, November 2008.

[Economics of the Singularity](#), *IEEE Spectrum*, 37-42, June 2008. Reprinted as: A New World Order, *Cosmos* 26:47-54, April/May 2009.

[Cut Medicine In Half](#), *CATO Unbound*, September 10, 2007. Replies by David Cutler, Dana Goldman, Alan Garber. Responses [Let Go The Monkey Trap](#), [Are the Aggregate Studies Misleading? Why?](#), [Still Seeking Specific Critiques](#), [Yes. There are Costs of Change](#), [Beware Double Standards](#), September 18,21,22,23,26.

[Birds of a Feather; Letter On Why Hawks Win](#), *Foreign Policy*, pp. 10-12, March/April 2007.

[You're Fired](#), *Forbes*, p.38, October 30, 2006.

[The Myth of Creativity](#). *Business Week* p.134, July 3, 2006.

[Reality and Fantasy in Economic Revolutions](#), *CATO Unbound*, Future of Work Issue, June 6, 2006. [mirror](#)

[Letter](#), *Wilson Quarterly* 30(2):7,9, Spring 2006.

[The Next Really Big Enormous Thing](#). *Future Brief*. October 6, 2004.

[The Morality of Meat - Why eating your vegetables is cruel to animals](#). *The American Enterprise Online Hot Flash*, July 24, 2002.

[Dreams of Autarky](#). *Doctor Tandy's Guide to Life Extension and Transhumanity*, ed. Charles Tandy, Ria University Press, 2001.

[Why Cryonics Isn't Popular](#). *Longevity Report* 12(72), July 1999.

Market-based approaches to managing science return from planetary missions, *European Space Agency (Special Publication) ESA SP 394(1):524-531*, 1996.

24 short essays in [DIY Futures - People's Ideas & Projects for a Better World](#), ed. [Nicholas Albery](#). [The Institute for Social Inventions](#), London, 1996. Also 4 short essays in *Creative Speculations*, 1997.

[Hard to Tax Scenario](#), with Hal Finney, Hettinga's Best of the Month, *Journal of Internet Banking and Finance* 2(1), January 1997.

[Idea Futures](#). *Wired* 3(9):125. [Excerpt](#) was Quotable Quote, *Wall Street Journal* A14, 30 August, 1995.

Links at <http://hanson.gmu.edu/vita.html>