Contemporary Statistics
Glamour Risk and Aftermath

D A S Fraser
Statistical Sciences U Toronto
Saw Swee Hock Visiting Prof U Hong Kong

Saw Swee Hock Public Lecture in Statistics
Dept Statistics & Actuarial Science U Hong Kong
Hong Kong Aug 22 2013

www.utstat.toronto.edu/dfraser/documents/266copss.pdf
/HK-I-Aug22.pdf

Statistics:

1	Data	Science	2011
2	Clinical trials	Science	2011
3	Vioxx	General	1999 +
4	L'Aquila	general	2009 +
5	Replication	Science	2011
6	Bayes goes to Washington	Science	2013
7	Bayes in 21st Century	Science	2013

- Directions
- Summary

Stalistics :

1	Data	Science	2011
2	Clinical trials	Science	2011
3	Vioxx	General	1999 +
4	L'Aquila	general	2009 +
5	Replication	Science	2011
6	Bayes goes to Washington	Science	2013
7	Bayes in 21st Century	Science	2013

- Directions
- Summary

Power, Risks, Responsibilities, Challenges

1 Science 2011 Feb11

Data

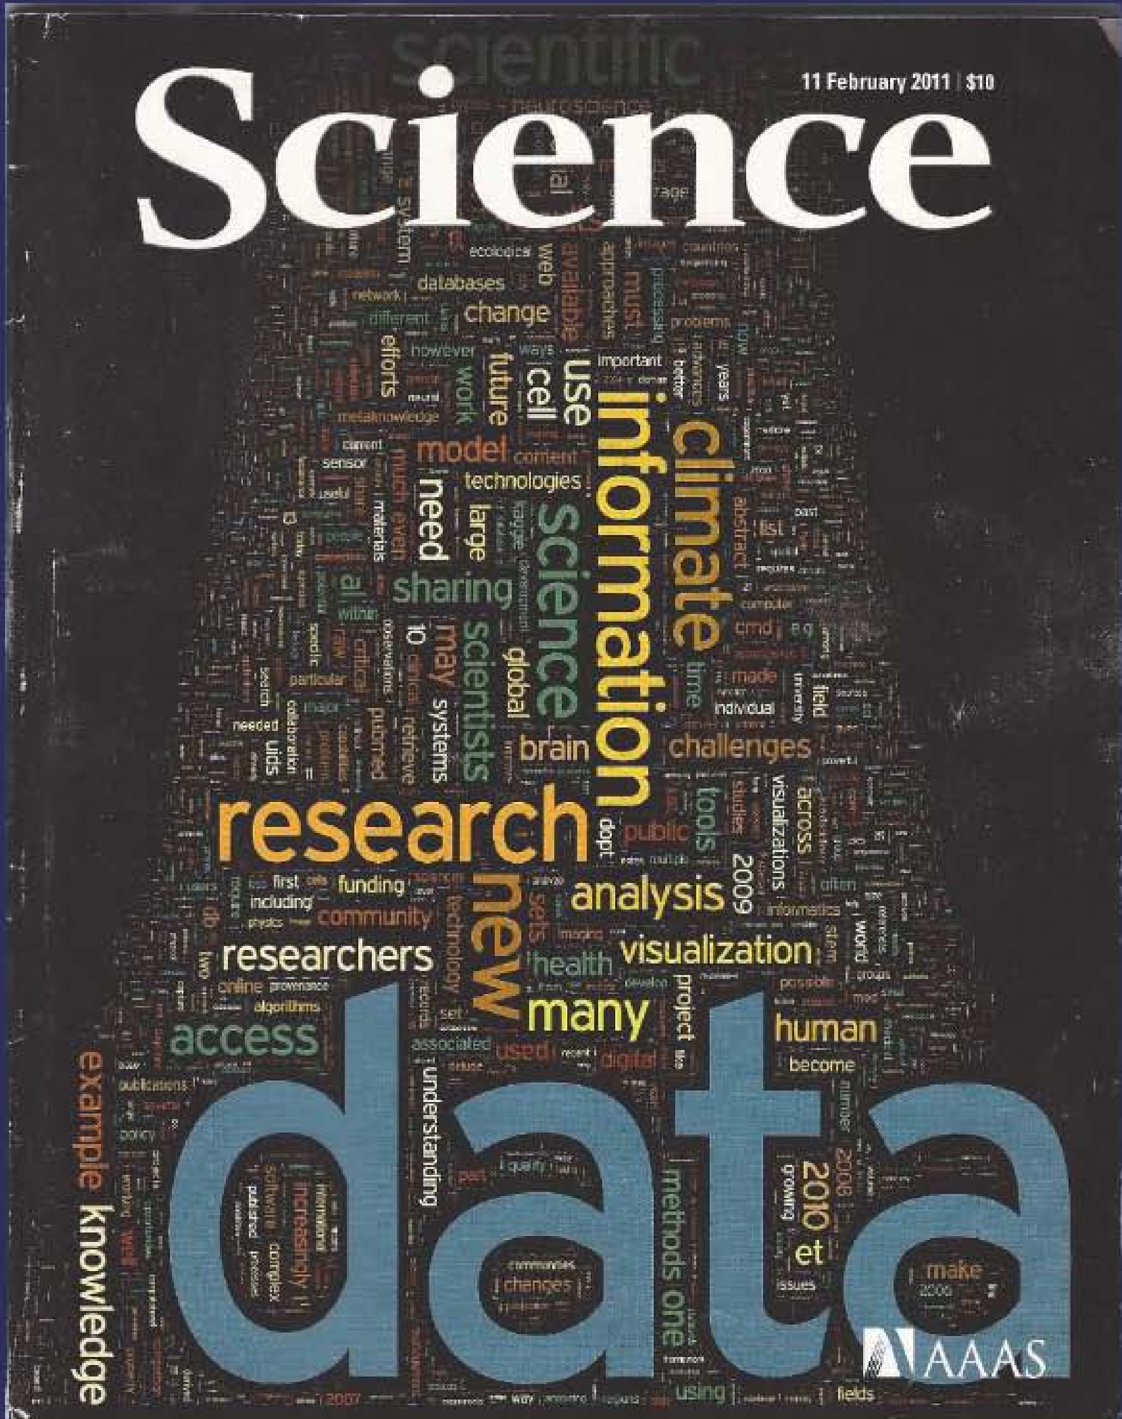

Science 2011 Feb 11

- a major science journal

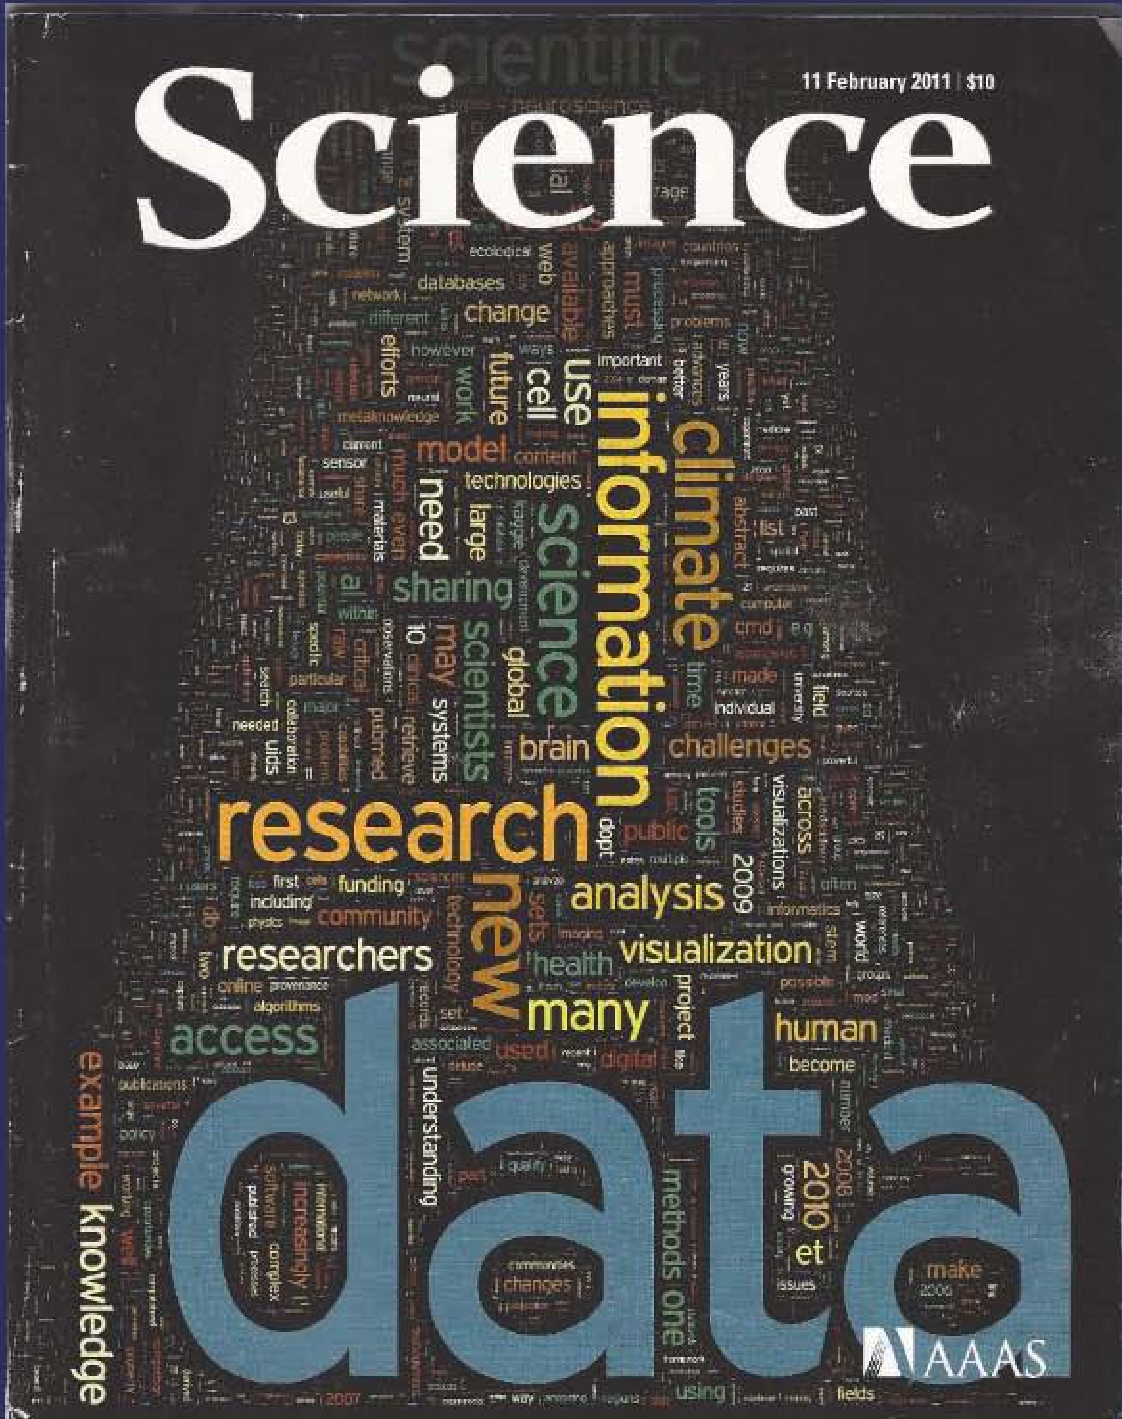

Science 2011 Feb 11

- a major science journal
- recognizes that

Data are everywhere!

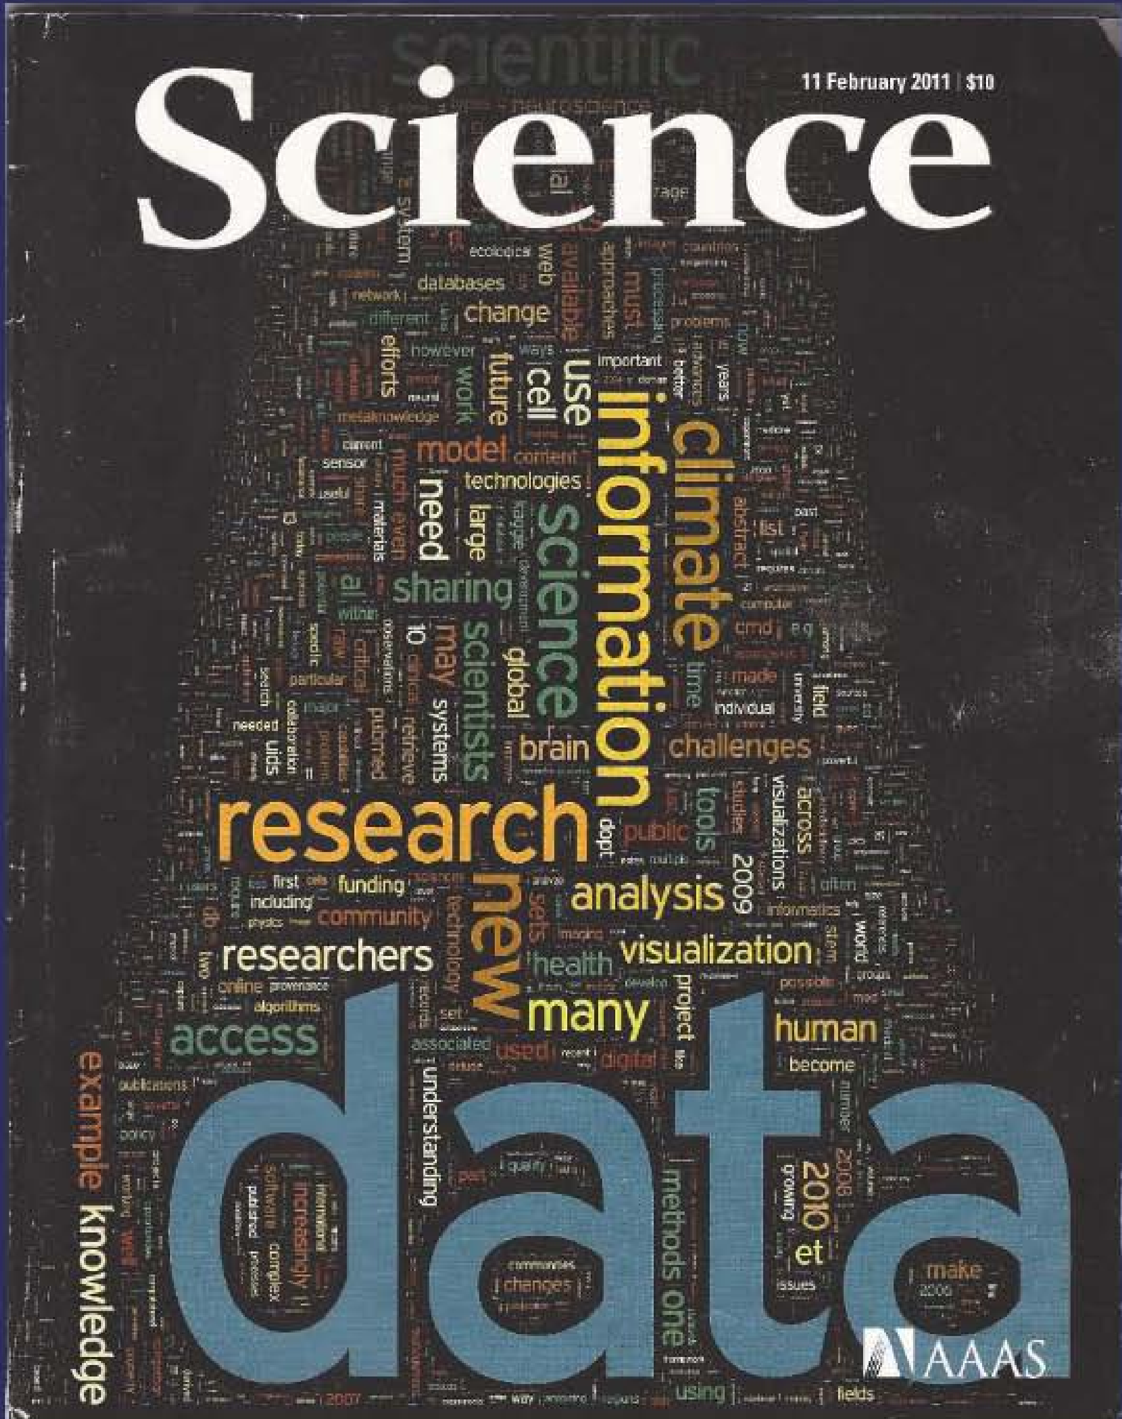

Science 2011 Feb 11

- a major science journal
- recognizes that

Data are everywhere!

This issue has

- 38 pages on Data
- 15 articles

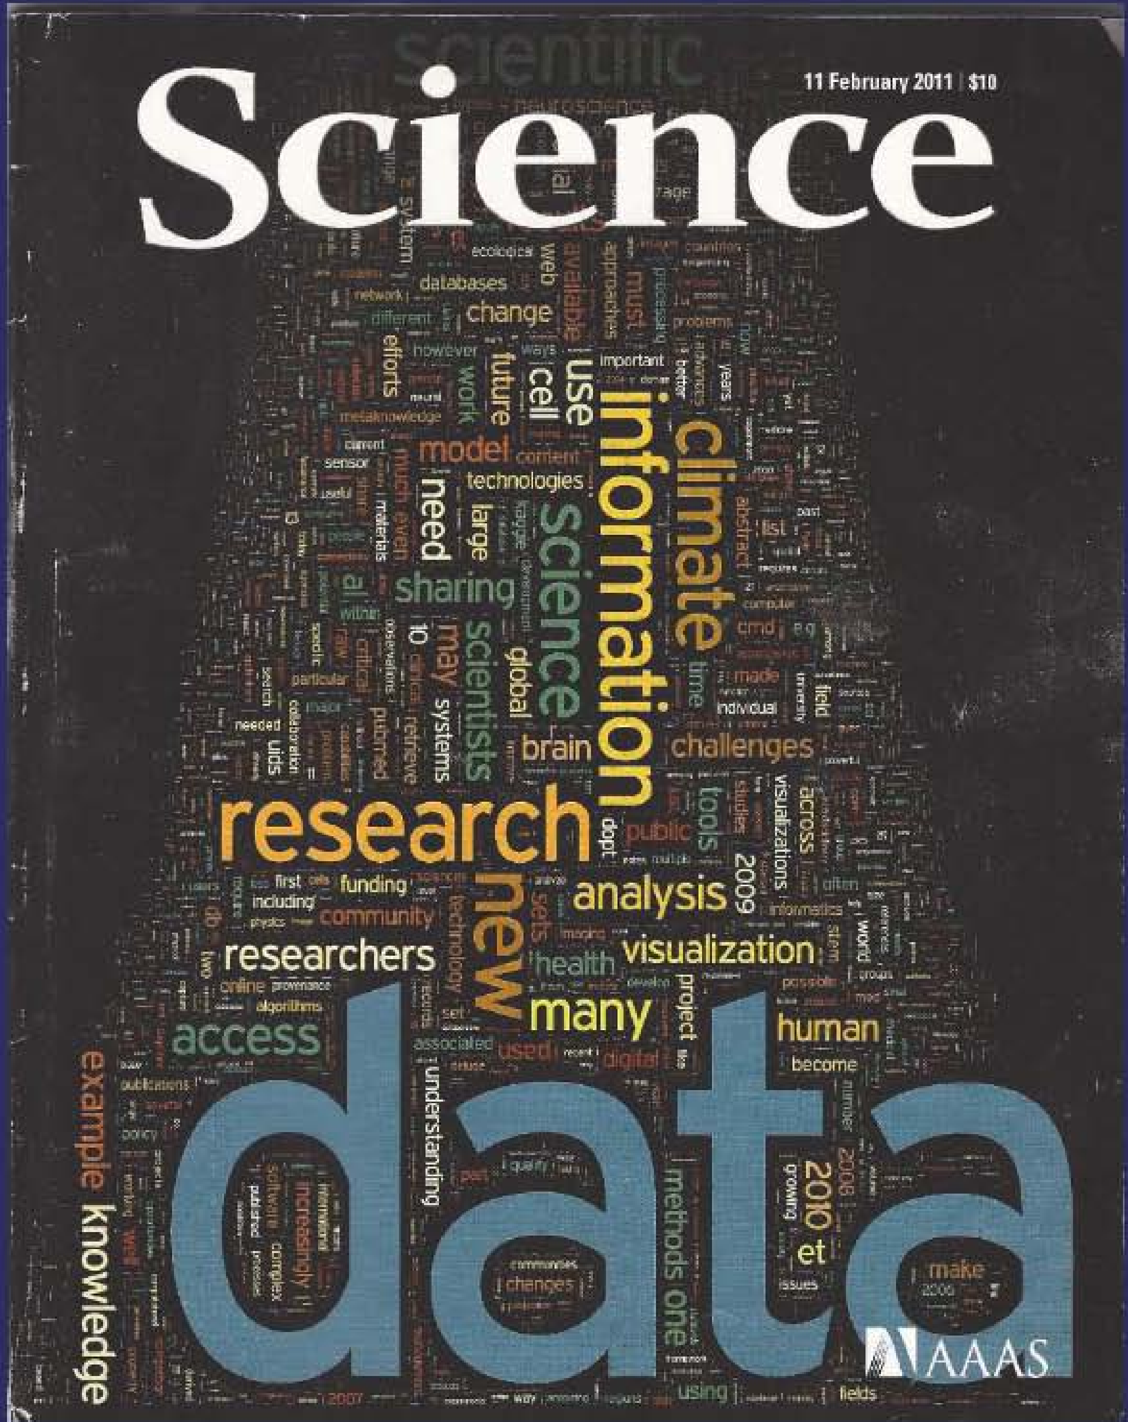

Science 2011 Feb 11

- a major science journal
- recognizes that

Data are everywhere!

This issue has

- 38 pages on Data
- 15 articles

See a "word cloud" <u>histogram</u>

- frequency of each word

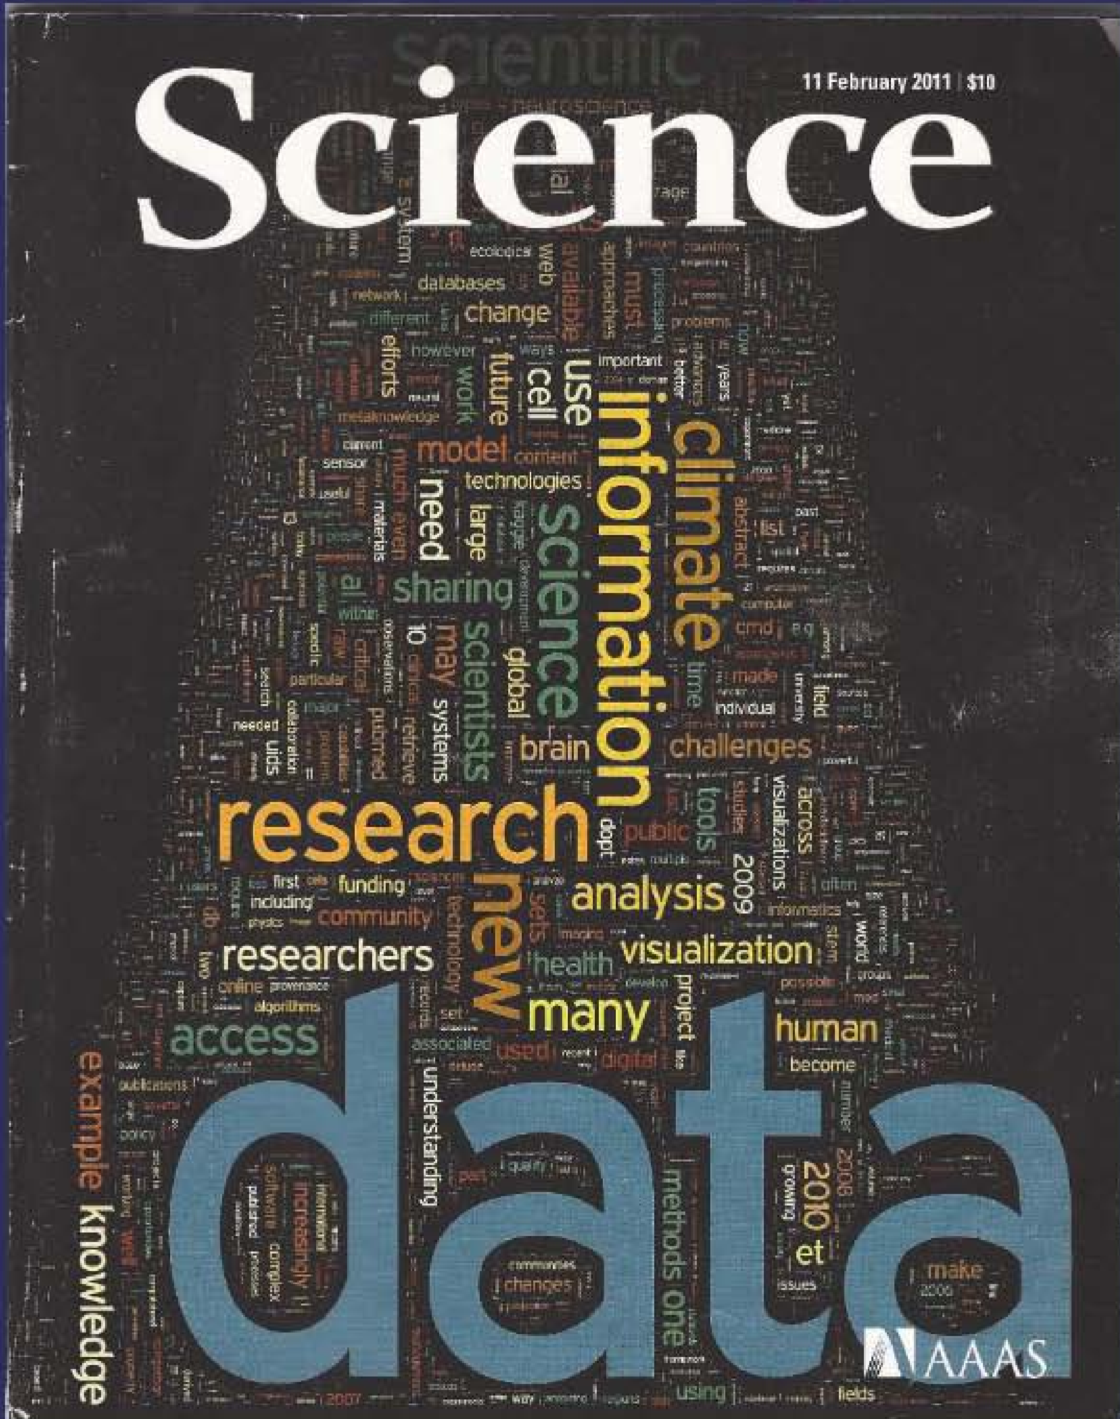

Science 2011 Feb 11

- a major science journal
- recognizes that

Data are everywhere!

This issue has

- 38 pages on Data
- 15 articles

See a "word cloud" <u>histogram</u>

- frequency of each word

But where is <u>statistics</u>?

"Not" there! (almost)

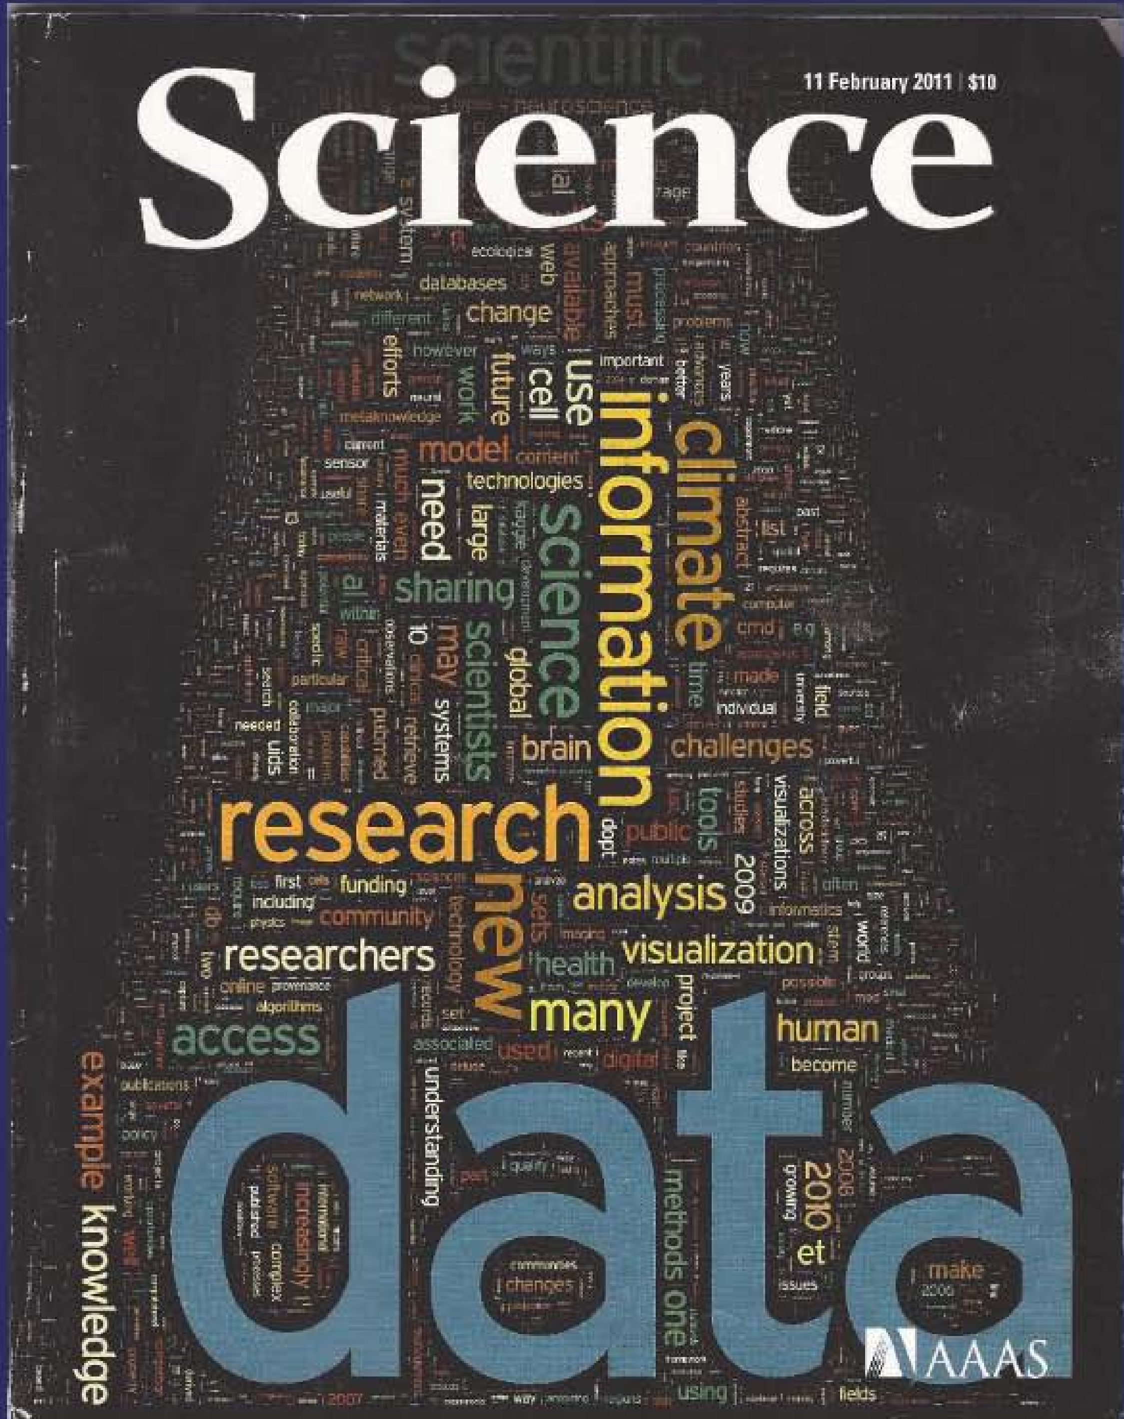

Science 2011 Feb 11

- a major science journal finds

No "Statistics"

in "Data" !

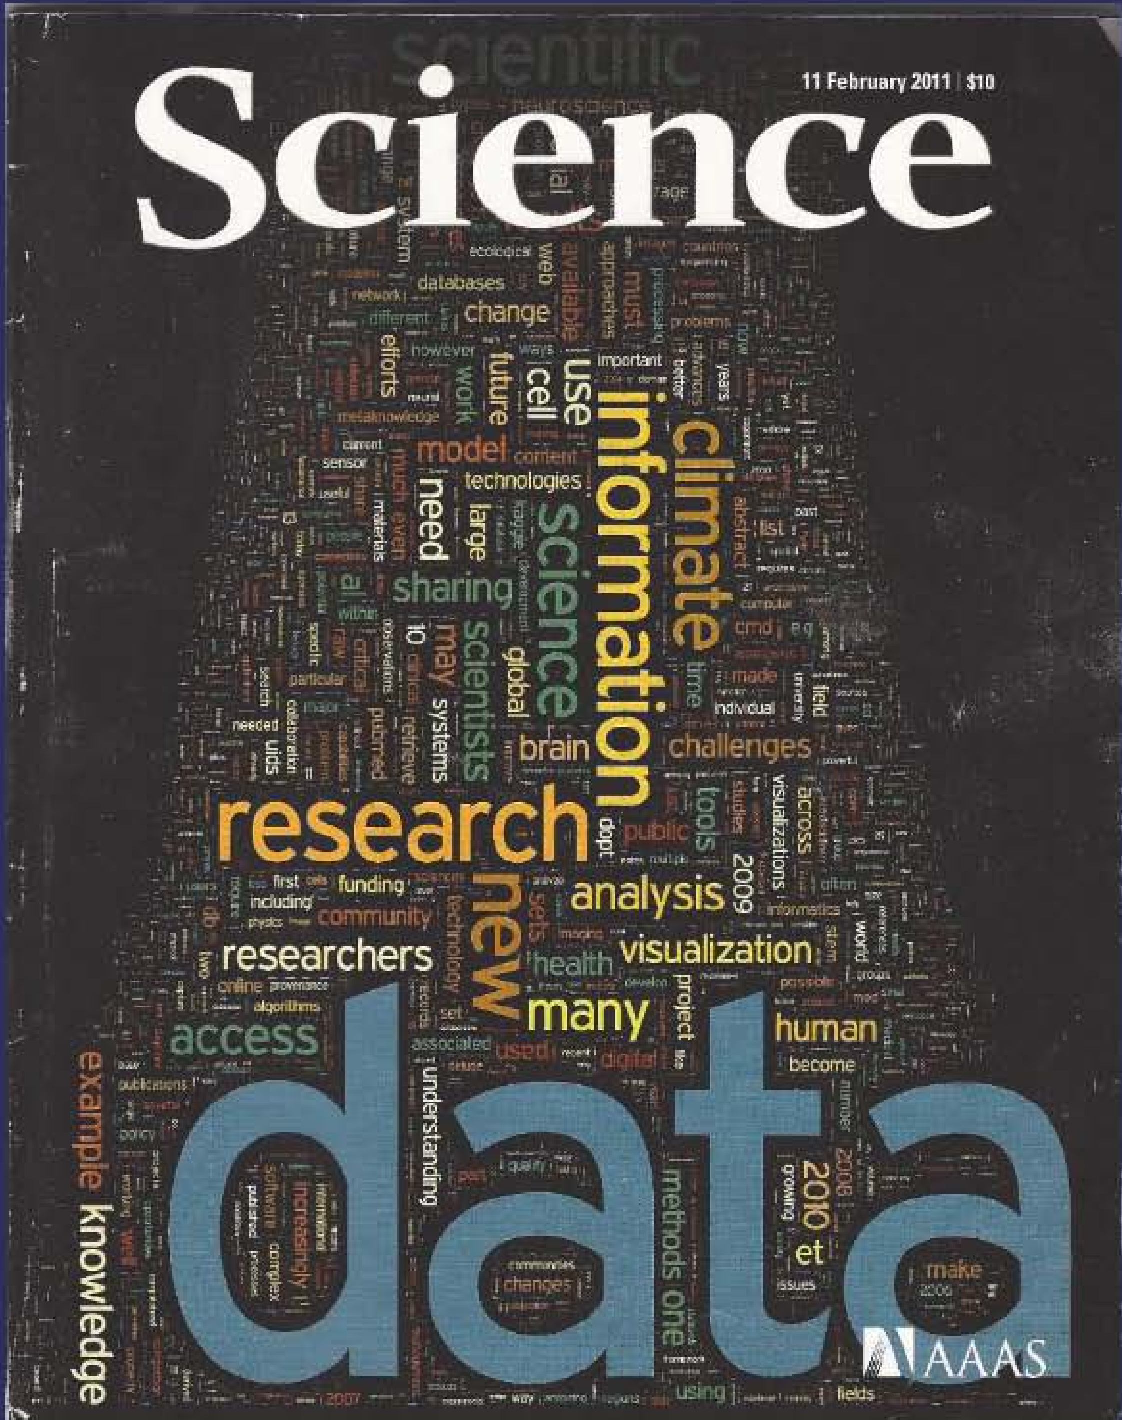

Science 2011 Feb 11

– a major science journal finds

No "Statistics" in "Data" !

"Statistics" is not part of Data ?

actually ...

2 Science 2011 Sep 23

Clinical Trials

Rethinking Clinical Trials

Andrew Grove is the former Chief Executive Officer of Intel Corporation. He is a patient advocate at the University of California, San Francisco.

THE BIOMEDICAL INDUSTRY SPENDS OVER $50 BILLION PER YEAR ON RESEARCH AND DEVELOPMENT and produces some 20 new drugs. One reason for this disappointing output is the byzantine U.S. clinical trial system that requires large numbers of patients. Half of all trials are delayed, 80 to 90% of them because of a shortage of trial participants. Patient limitations also cause large and unpredicted expenses to pharmaceutical and biotech companies as they are forced to tread water. As the industry moves toward biologics and personalized medicine, this limitation will become even greater. A breakthrough in regulation is needed to create a system that does more with fewer patients.

The current clinical trial system in the United States is more than 50 years old. Its architecture was conceived when electronic manipulation of data was limited, slow, and expensive. Since then, network and connectivity costs have declined ten thousand–fold, data storage costs over a million-fold, and computation costs by an even larger factor. Today, complex and powerful applications like electronic commerce are deployed on a large scale. Amazon.com is a good example. A large database of customers and products form the kernel of its operation. A customer's characteristics (like buying history and preferences) are observed and stored. Customers can be grouped and the buying behavior of any individual or group can be compared with corresponding behavior of others. Amazon can also track how a group or an individual responds to an outside action (such as advertising).

We might conceptualize an "e-trial" system along similar lines. Drug safety would continue to be ensured by the U.S. Food and Drug Administration. While safety-focused Phase I trials would continue under their jurisdiction, establishing efficacy would no longer be under their purview. Once safety is proven, patients could access the medicine in question through qualified physicians. Patients' responses to a drug would be stored in a database, along with their medical histories. Patient identity would be protected by biometric identifiers, and the database would be open to qualified medical researchers as a "commons." The response of any patient or group of patients to a drug or treatment would be tracked and compared to those of others in the database who were treated in a different manner or not at all. These comparisons would provide insights into the factors that determine real-life efficacy: how individuals or subgroups respond to the drug. This would liberate drugs from the tyranny of the averages that characterize trial information today. The technology would facilitate such comparisons at incredible speeds and could quickly highlight negative results. As the patient population in the database grows and time passes, analysis of the data would also provide the information needed to conduct postmarketing studies and comparative effectiveness research.

Today's e-commerce systems started small and took nearly 20 years to develop. Adapting this kind of capability to medical information would be a monumental undertaking. Initiating and overseeing it would be an appropriate task for the professional societies. There are encouraging signs, including a call in 2004 by the American Medical Association for public registries of drugs, as well as a proposal for trials that incorporate feed-forward mechanisms.* Another proposal would allow patients to choose between medicines whose efficacy has been determined in different manners.† There is also a suggestion to use control of pricing to encourage drug developers to move forward in a "progressive" trial design.‡ Ideas, however, are not enough. We need the professions to mobilize and take advantage of this enormous opportunity.

– Andrew Grove

10.1126/science.1212118

*H. Hvitfeldt et al., Qual. Manag. Health Care 18, 247 (2009). †B. J. Madden, Free to Choose Medicine (Heartland Institute, Chicago, IL, 2010). ‡M. Boldrin, S. J. Swamidass, "A New Bargain for Drug Approvals," Wall Street Journal, 25 July 2011.

Science 2011 Sep 23

Rethinking Clinical Trials

Andrew Grove is the former Chief Executive Officer of Intel Corporation. He is a patient advocate at the University of California, San Francisco.

THE BIOMEDICAL INDUSTRY SPENDS OVER $50 BILLION PER YEAR ON RESEARCH AND DEVELOPMENT and produces some 20 new drugs. One reason for this disappointing output is the byzantine U.S. clinical trial system that requires large numbers of patients. Half of all trials are delayed, 80 to 90% of them because of a shortage of trial participants. Patient limitations also cause large and unpredicted expenses to pharmaceutical and biotech companies as they are forced to tread water. As the industry moves toward biologics and personalized medicine, this limitation will become even greater. A breakthrough in regulation is needed to create a system that does more with fewer patients.

The current clinical trial system in the United States is more than 50 years old. Its architecture was conceived when electronic manipulation of data was limited, slow, and expensive. Since then, network and connectivity costs have declined ten thousand–fold, data storage costs over a million-fold, and computation costs by an even larger factor. Today, complex and powerful applications like electronic commerce are deployed on a large scale. Amazon.com is a good example. A large database of customers and products form the kernel of its operation. A customer's characteristics (like buying history and preferences) are observed and stored. Customers can be grouped and the buying behavior of any individual or group can be compared with corresponding behavior of others. Amazon can also track how a group or an individual responds to an outside action (such as advertising).

We might conceptualize an "e-trial" system along similar lines. Drug safety would continue to be ensured by the U.S. Food and Drug Administration. While safety-focused Phase I trials would continue under their jurisdiction, establishing efficacy would no longer be under their purview. Once safety is proven, patients could access the medicine in question through qualified physicians. Patients' responses to a drug would be stored in a database, along with their medical histories. Patient identity would be protected by biometric identifiers, and the database would be open to qualified medical researchers as a "commons." The response of any patient or group of patients to a drug or treatment would be tracked and compared to those of others in the database who were treated in a different manner or not at all. These comparisons would provide insights into the factors that determine real-life efficacy: how individuals or subgroups respond to the drug. This would liberate drugs from the tyranny of the averages that characterize trial information today. The technology would facilitate such comparisons at incredible speeds and could quickly highlight negative results. As the patient population in the database grows and time passes, analysis of the data would also provide the information needed to conduct postmarketing studies and comparative effectiveness research.

Today's e-commerce systems started small and took nearly 20 years to develop. Adapting this kind of capability to medical information would be a monumental undertaking. Initiating and overseeing it would be an appropriate task for the professional societies. There are encouraging signs, including a call in 2004 by the American Medical Association for public registries of drugs, as well as a proposal for trials that incorporate feed-forward mechanisms.* Another proposal would allow patients to choose between medicines whose efficacy has been determined in different manners.† There is also a suggestion to use control of pricing to encourage drug developers to move forward in a "progressive" trial design.‡ Ideas, however, are not enough. We need the professions to mobilize and take advantage of this enormous opportunity.

– Andrew Grove

10.1126/science.1212118

*H. Hvitfeldt et al., *Qual. Manag. Health Care* 18, 247 (2009). †B. J. Madden, *Free to Choose Medicine* (Heartland Institute, Chicago, IL, 2010). ‡M. Boldrin, S. J. Swamidass, "A New Bargain for Drug Approvals," *Wall Street Journal*, 25 July 2011.

Science 2011 Sep 23

Grove: former CEO Intel

Andrew Grove is the former Chief Executive Officer of Intel Corporation. He is a patient advocate at the University of California, San Francisco.

Rethinking Clinical Trials

THE BIOMEDICAL INDUSTRY SPENDS OVER $50 BILLION PER YEAR ON RESEARCH AND DEVELOPMENT and produces some 20 new drugs. One reason for this disappointing output is the byzantine U.S. clinical trial system that requires large numbers of patients. Half of all trials are delayed, 80 to 90% of them because of a shortage of trial participants. Patient limitations also cause large and unpredicted expenses to pharmaceutical and biotech companies as they are forced to tread water. As the industry moves toward biologics and personalized medicine, this limitation will become even greater. A breakthrough in regulation is needed to create a system that does more with fewer patients.

The current clinical trial system in the United States is more than 50 years old. Its architecture was conceived when electronic manipulation of data was limited, slow, and expensive. Since then, network and connectivity costs have declined ten thousand–fold, data storage costs over a million-fold, and computation costs by an even larger factor. Today, complex and powerful applications like electronic commerce are deployed on a large scale. Amazon.com is a good example. A large database of customers and products form the kernel of its operation. A customer's characteristics (like buying history and preferences) are observed and stored. Customers can be grouped and the buying behavior of any individual or group can be compared with corresponding behavior of others. Amazon can also track how a group or an individual responds to an outside action (such as advertising).

We might conceptualize an "e-trial" system along similar lines. Drug safety would continue to be ensured by the U.S. Food and Drug Administration. While safety-focused Phase I trials would continue under their jurisdiction, establishing efficacy would no longer be under their purview. Once safety is proven, patients could access the medicine in question through qualified physicians. Patients' responses to a drug would be stored in a database, along with their medical histories. Patient identity would be protected by biometric identifiers, and the database would be open to qualified medical researchers as a "commons." The response of any patient or group of patients to a drug or treatment would be tracked and compared to those of others in the database who were treated in a different manner or not at all. These comparisons would provide insights into the factors that determine real-life efficacy: how individuals or subgroups respond to the drug. This would liberate drugs from the tyranny of the averages that characterize trial information today. The technology would facilitate such comparisons at incredible speeds and could quickly highlight negative results. As the patient population in the database grows and time passes, analysis of the data would also provide the information needed to conduct postmarketing studies and comparative effectiveness research.

Today's e-commerce systems started small and took nearly 20 years to develop. Adapting this kind of capability to medical information would be a monumental undertaking. Initiating and overseeing it would be an appropriate task for the professional societies. There are encouraging signs, including a call in 2004 by the American Medical Association for public registries of drugs, as well as a proposal for trials that incorporate feed-forward mechanisms.* Another proposal would allow patients to choose between medicines whose efficacy has been determined in different manners.† There is also a suggestion to use control of pricing to encourage drug developers to move forward in a "progressive" trial design.‡ Ideas, however, are not enough. We need the professions to mobilize and take advantage of this enormous opportunity.

– Andrew Grove

10.1126/science.1212118

*H. Hvitfeldt et al., *Qual. Manag. Health Care* 18, 247 (2009). †B. J. Madden, *Free to Choose Medicine* (Heartland Institute, Chicago, IL, 2010). ‡M. Boldrin, S. J. Swamidass, "A New Bargain for Drug Approvals," *Wall Street Journal*, 25 July 2011.

Science 2011 Sep 23

Grove: former CEO Intel
- clinical trials out-dated
- Large data base - customers
  - products
- open to researchers

and

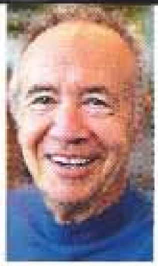

Andrew Grove is the former Chief Executive Officer of Intel Corporation. He is a patient advocate at the University of California, San Francisco.

EDITORIAL

Rethinking Clinical Trials

THE BIOMEDICAL INDUSTRY SPENDS OVER $50 BILLION PER YEAR ON RESEARCH AND DEVELOPMENT and produces some 20 new drugs. One reason for this disappointing output is the byzantine U.S. clinical trial system that requires large numbers of patients. Half of all trials are delayed, 80 to 90% of them because of a shortage of trial participants. Patient limitations also cause large and unpredicted expenses to pharmaceutical and biotech companies as they are forced to tread water. As the industry moves toward biologics and personalized medicine, this limitation will become even greater. A breakthrough in regulation is needed to create a system that does more with fewer patients.

The current clinical trial system in the United States is more than 50 years old. Its architecture was conceived when electronic manipulation of data was limited, slow, and expensive. Since then, network and connectivity costs have declined ten thousand–fold, data storage costs over a million-fold, and computation costs by an even larger factor. Today, complex and powerful applications like electronic commerce are deployed on a large scale. Amazon.com is a good example. A large database of customers and products form the kernel of its operation. A customer's characteristics (like buying history and preferences) are observed and stored. Customers can be grouped and the buying behavior of any individual or group can be compared with corresponding behavior of others. Amazon can also track how a group or an individual responds to an outside action (such as advertising).

We might conceptualize an "e-trial" system along similar lines. Drug safety would continue to be ensured by the U.S. Food and Drug Administration. While safety-focused Phase I trials would continue under their jurisdiction, establishing efficacy would no longer be under their purview. Once safety is proven, patients could access the medicine in question through qualified physicians. Patients' responses to a drug would be stored in a database, along with their medical histories. Patient identity would be protected by biometric identifiers, and the database would be open to qualified medical researchers as a "commons." The response of any patient or group of patients to a drug or treatment would be tracked and compared to those of others in the database who were treated in a different manner or not at all. These comparisons would p… real-life efficacy: how individuals or subgro… drugs from the tyranny of the averages that c… nology would facilitate such comparisons at incredible speeds and could quickly highlight negative results. As the patient population in the database grows and time passes, analysis of the data would also provide the information needed to conduct postmarketing studies and comparative effectiveness research.

Today's e-commerce systems started small and took nearly 20 years to develop. Adapting this kind of capability to medical information would be a monumental undertaking. Initiating and overseeing it would be an appropriate task for the professional societies. There are encouraging signs, including a call in 2004 by the American Medical Association for public registries of drugs, as well as a proposal for trials that incorporate feed-forward mechanisms.* Another proposal would allow patients to choose between medicines whose efficacy has been determined in different manners.† There is also a suggestion to use control of pricing to encourage drug developers to move forward in a "progressive" trial design.‡ Ideas, however, are not enough. We need the professions to mobilize and take advantage of this enormous opportunity.

– Andrew Grove

10.1126/science.1212118

*H. Hvitfeldt *et al.*, *Qual. Manag. Health Care* 18, 247 (2009). †R. J. Madden, *Free to Choose Medicine* (Heartland Institute, Chicago, IL, 2010). ‡M. Boldrin, S. J. Swamidass, "A New Bargain for Drug Approvals," *Wall Street Journal*, 25 July 2011.

CREDITS: (LEFT) MICHAEL PEYCE PHOTOGRAPHY; (RIGHT) IMAGES.COM/CORBIS

Science 2011 Sep 23

Grove: former CEO Intel
- clinical trials out-dated
- Large data base - customers
  - products
- open to researchers

and

Once safety is proven, patients could access the medicine in question through qualified physicians.

Andrew Grove is the former Chief Executive Officer of Intel Corporation. He is a patient advocate at the University of California, San Francisco.

Rethinking Clinical Trials

THE BIOMEDICAL INDUSTRY SPENDS OVER $50 BILLION PER YEAR ON RESEARCH AND DEVELOPMENT and produces some 20 new drugs. One reason for this disappointing output is the byzantine U.S. clinical trial system that requires large numbers of patients. Half of all trials are delayed, 80 to 90% of them because of a shortage of trial participants. Patient limitations also cause large and unpredicted expenses to pharmaceutical and biotech companies as they are forced to tread water. As the industry moves toward biologics and personalized medicine, this limitation will become even greater. A breakthrough in regulation is needed to create a system that does more with fewer patients.

The current clinical trial system in the United States is more than 50 years old. Its architecture was conceived when electronic manipulation of data was limited, slow, and expensive. Since then, network and connectivity costs have declined ten thousand–fold, data storage costs over a million-fold, and computation costs by an even larger factor. Today, complex and powerful applications like electronic commerce are deployed on a large scale. Amazon.com is a good example. A large database of customers and products form the kernel of its operation. A customer's characteristics (like buying history and preferences) are observed and stored. Customers can be grouped and the buying behavior of any individual or group can be compared with corresponding behavior of others. Amazon can also track how a group or an individual responds to an outside action (such as advertising).

We might conceptualize an "e-trial" system along similar lines. Drug safety would continue to be ensured by the U.S. Food and Drug Administration. While safety-focused Phase I trials would continue under their jurisdiction, establishing efficacy would no longer be under their purview. Once safety is proven, patients could access the medicine in question through qualified physicians. Patients' responses to a drug would be stored in a database, along with their medical histories. Patient identity would be protected by biometric identifiers, and the database would be open to qualified medical researchers as a "commons." The response of any patient or group of patients to a drug or treatment would be tracked and compared to those of others in the database who were treated in a different manner or not at all. These comparisons would [illegible] real-life efficacy: how individuals or subgro[illegible] drugs from the tyranny of the averages that [illegible] nology would facilitate such comparisons at incredible speeds and could quickly highlight negative results. As the patient population in the database grows and time passes, analysis of the data would also provide the information needed to conduct postmarketing studies and comparative effectiveness research.

Today's e-commerce systems started small and took nearly 20 years to develop. Adapting this kind of capability to medical information would be a monumental undertaking. Initiating and overseeing it would be an appropriate task for the professional societies. There are encouraging signs, including a call in 2004 by the American Medical Association for public registries of drugs, as well as a proposal for trials that incorporate feed-forward mechanisms.* Another proposal would allow patients to choose between medicines whose efficacy has been determined in different manners.† There is also a suggestion to use control of pricing to encourage drug developers to move forward in a "progressive" trial design.‡ Ideas, however, are not enough. We need the professions to mobilize and take advantage of this enormous opportunity.

– Andrew Grove

10.1126/science.1212118

*H. Hvitfeldt et al., Qual. Manag. Health Care 18, 247 (2009). †B. J. Madden, Free to Choose Medicine (Heartland Institute, Chicago, IL, 2010). ‡M. Boldrin, S. J. Swamidass, "A New Bargain for Drug Approvals," Wall Street Journal, 25 July 2011.

Science 2011 Sep 23

Grove: former CEO Intel

- clinical trials out-dated
- Large data base - customers
  - products
- open to researchers

and

Once safety is proven, patients could access the medicine in question through qualified physicians.

Once safety is proven,

is <u>not</u> so simple

"Vioxx ....."

3 Vioxx

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

1999 Approved by FDA (8 year approval process)
2000 Advertizing \$160m .... More than Pepsi or Budweiser
2003 Earned \$2.46 b Big bucks!

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

  1999 Approved by FDA (8 year approval process)

  2000 Advertizing $160m .... More than Pepsi or Budweiser

  2003 Earned $2.46 b Big bucks!

- Evidence available: of 3x, 5x rate of CVT events (cardio-vascular T)

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

  1999 Approved by FDA (8 year approval process)

  2000 Advertizing \$160m .... More than Pepsi or Budweiser

  2003 Earned \$2.46 b Big bucks!

- Evidence available: of 3x, 5x rate of CV events (cardio-vascular ↑)

- litigation vs. Merck: 2004 ... by those affected

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"
  - 1999 Approved by FDA (8 year approval process)
  - 2000 Advertizing \$160m .... More than Pepsi or Budweiser
  - 2003 Earned \$2.46 b Big bucks!
- Evidence available: of 3x, 5x rate of CVT events (cardio-vascular T)
- litigation vs. Merck: 2004
  - that M knew of increased incidence of CV events

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

  1999 Approved by FDA (8 year approval process)

  2000 Advertizing $160m .... More than Pepsi or Budweiser

  2003 Earned $2.46 b Big bucks!

- Evidence available: of 3x, 5x rate of CVT events (cardio-vascular T)

- litigation vs. Merck: 2004

  - that M knew of increased incidence of CV events

  - that M had "filtered" Data

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"
  - 1999 Approved by FDA (8 year approval process)
  - 2000 Advertizing \$160m .... More than Pepsi or Budweiser
  - 2003 Earned \$2.46 b Big bucks!
- Evidence available: of 3x, 5x rate of CV events (cardio-vascular T)
- litigation vs. Merck: 2004
  - that M knew of increased incidence of CV events
  - that M had "filtered" Data
  - that V linked to 30,000 to 40,000 deaths

Vioxx

- An analgesic (pain relief) from Merck "Pharma"

  1999 Approved by FDA (8 year approval process)

  2000 Advertizing $160m .... More than Pepsi or Budweiser

  2003 Earned $2.46 b Big bucks !

- Evidence available: of 3x, 5x rate of CV events (cardio-vascular ↑)

- litigation vs. Merck: 2004

  - that M knew of increased incidence of CV events
  - that M had "filtered" Data
  - that V linked to 30,000 to 40,000 deaths
  - David Madigan, Chair, Columbia U Statistics (litigation consultant)

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

1999 Approved by FDA
2000 Advertizing \$160m .... More than Pepsi or Budweiser
2003 Earned \$2.46 b Big bucks!

- "Proved" safe
- 30,000 - 40,000 deaths

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

1999 Approved by FDA

2000 Advertizing \$160m .... More than Pepsi or Budweiser

2003 Earned \$2.46 b Big bucks!

- "Proved" safe
- 30,000 - 40,000 deaths
- \$5b penalty

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

1999 Approved by FDA

2000 Advertizing \$160m .... More than Pepsi or Budweiser

2003 Earned \$2.46 b Big bucks!

- "Proved" safe
- 30,000 - 40,000 deaths
- \$5b penalty

"Merck got a bargain" cf profit

3 Vioxx

- An analgesic (pain relief) from Merck "Pharma"

1999 Approved by FDA
2000 Advertizing $160m .... More than Pepsi or Budweiser
2003 Earned $2.46 b Big bucks!

- "Proved" safe
- 30,000 - 40,000 deaths
- $5b penalty

"Merck got a bargain"

Message ?

4 L'Aquila 2009 Apr 6

4 L'Aquila 2009 Apr 6 Significance 2012 Dec

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

L'Aquila 2009 Apr 6 Significance 2012 Dec

feature

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Firemen touch a marble statue of the Madonna after removing it from the top of the church in Paganica, near L'Aquila. Photo: Reuters

firemen removing
Madonna

L'Aquila 2009 Apr 6 Significance 2012 Dec

Earthquake 2009 Apr 6
300 deaths

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages; there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

L'Aquila 2009 Apr 6 Significance 2012 Dec

feature

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages; there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock

L'Aquila 2009 Apr 6 Significance 2012 Dec

feature

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock
Panicked residents

L'Aquila 2009 Apr 6 Significance 2012 Dec

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock
Panicked residents
Committee of seismologists formed

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

L'Aquila 2009 Apr 6 Significance 2012 Dec

feature

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock
Panicked residents

Committee of seismologists formed

Reported: no reason to expect a big quake
- 'small shocks reducing stresses'

L'Aquila 2009 Apr 6 Significance 2012 Dec

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock
Panicked residents

Committee of seismologists formed

Reported: no reason to expect a big quake
- 'small shocks reducing stresses'

but earthquake came...
300 died'

L'Aquila 2009 Apr 6 Significance 2012 Dec

feature

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Gianmichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock
Panicked residents

Committee of seismologists formed

Reported: no reason to expect
a big quake
- 'small shocks reducing stresses'

Courts convicted 7 of manslaughter

L'Aquila 2009 Apr 6 Significance 2012 Dec

The L'Aquila earthquake

Science or risk on trial?

On April 6th 2009 a major earthquake devastated the Italian town of L'Aquila, with the loss of 309 lives. In October this year six scientists and a government official were sentenced to 6 years each for providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and for falsely reassuring the population. **Jordi Prats** examines the case.

The L'Aquila earthquake was such a shocking, calamitous event that, as with 9/11, Italians remember where they were when they heard of it. In 2009, I was living in Bari, an Italian city about 320 km to the south-east of L'Aquila, and a friend of mine was visiting me during her Easter holidays. During the night of April 6th she woke suddenly, feeling as if her bed had leapt. I am a heavy sleeper; I sensed nothing of the kind. Had she dreamt it? Unfortunately, she had not. In the morning the television was full of the news of the disaster. The earthquake had hit the city of L'Aquila and many surrounding villages: there were 309 deaths, more than 25 000 displaced people, and more than 10 000 significantly damaged buildings. Reconstruction costs were estimated at more than US$16 billion. In the days that followed, a wave of solidarity with Aquilanians and the people of the Abruzzo region spread over the country. But there was a widespread opinion that most of those lives could have been saved if only the Serious Risks Commission, a counselling panel advising the Italian Civil Protection Agency, had done a better job of evaluating and communicating seismic risk.

As a result, the members of the Serious Risk Commission were taken to court. Professors Franco Barberi, Enzo Boschi, Giammichele Calvi, Claudio Eva, Mauro Dolce and Giulio Selvaggi can be counted among the best geoscientists in Italy. With them was Bernardo De Bernardinis, a government administrator and former vice-president of the Italian Civil Protection Agency's technical department. They were accused of providing "inaccurate, incomplete and contradictory information" on the probability and risk of an earthquake and of falsely reassuring the population. Their accusers argued that the soothing declarations made by the commission made people change their behaviour in the face of the quake. Confident in scientific advice, they chose to stay at home instead of fleeing when the earthquake began, and were caught by the collapsing buildings. The Italian scientists were thus held responsible for the deaths

Government offices, L'Aquila. Photo: Reuters

Earthquake 2009 Apr 6
300 deaths

But early 2009
Many small shock
Panicked residents

Committee of seismologists formed
Reported: no reason to expect a big quake
- 'small shocks reducing stresses'

Courts convicted 7 of manslaughter

Statistics was involved!

Where does responsibility fall?

5 Science 2011 Dec 2

Replication and Reproducibility

5

Science 2011 Dec 2

INTRODUCTION

Again, and Again, and Again ...

REPLICATION—THE CONFIRMATION OF RESULTS AND CONCLUSIONS FROM ONE STUDY obtained independently in another—is considered the scientific gold standard. New tools and technologies, massive amounts of data, long-term studies, interdisciplinary approaches, and the complexity of the questions being asked are complicating replication efforts, as are increased pressures on scientists to advance their research. The five Perspectives in this section (and associated News and Careers stories, Readers' Poll, and Editorial) explore some of the issues associated with replicating results across various fields.

Ryan (p. 1229) highlights the excitement and challenges that come with field-based research. In particular, observing processes as they occur in nature allows for discovery but makes replication difficult, because the precise conditions surrounding the observations are unique. Further, although laboratory research allows for the specification of experimental conditions, the conclusions may not apply to the real world. Debate about the merits of lab-based and field-based studies has been a persistent theme over time. Tomasello and Call (p. 1227) further contribute to this debate in their discussion of some obvious barriers to replication in primate cognition and behavior research (small numbers of subjects, expense, and ethics issues) as well as more subtle ones, such as the nontrivial challenge of designing tasks that elicit complex cognitive behaviors.

New technologies continue to produce a deluge of data of different varieties, raising expectations for new knowledge that will translate into meaningful therapeutics and insights into health. Ioannidis and Khoury (p. 1230) outline multiple steps for validating such large-scale data on the path to clinical utility and make suggestions for incentives (and penalties) that could enhance the availability of reliable data and analyses.

Peng (p. 1226) discusses the need for a minimum standard of reproducibility in computer sciences, arguing that enough information about methods and code should be available for independent researchers to reach consistent conclusions using original raw data. Specifically, he describes a model that one journal has used to make this a reality.

The need to convince the public that data are replicable has grown as science and public policy-making intersect, an issue that has has beset climate change studies. As Santer *et al.* (p. 1232) describe, having multiple groups examining the same data and generating new data has led to robust conclusions.

The importance of replication and reproducibility for scientists is unquestioned. Sometimes attempts to replicate reveal scientific uncertainties. This is one of the main ways that science progresses (see associated News stories of faster-than-light neutrinos and sirtuins, pp. 1200 and 1194). Unfortunately, in rare instances (compared to the body of scientific work), it can also indicate fraud (see the Editorial by Crocker and Cooper, p. 1182). How do we promote the publication of replicable data? The authors in this section come up with possibilities that are targeted at funders, journals, and the research culture itself. In the Readers' Poll, you can make your views known as well.

– BARBARA R. JASNY, GILBERT CHIN, LISA CHONG, SACHA VIGNIERI

Data Replication & Reproducibility

CONTENTS

Perspectives

See also Editorial p. 1182; News stories pp. 1194 and 1200; Readers' Poll p. 1202; Science Careers content p. 1179; and www.sciencemag.org/special/data-rep/

Science

INTRODUCTION

Again, and Again, and Again . . .

REPLICATION—THE CONFIRMATION OF RESULTS AND CONCLUSIONS FROM ONE STUDY obtained independently in another—is considered the scientific gold standard. New tools and technologies, massive amounts of data, long-term studies, interdisciplinary approaches, and the complexity of the questions being asked are complicating replication efforts, as are increased pressures on scientists to advance their research. The five Perspectives in this section (and associated News and Careers stories, Readers' Poll, and Editorial) explore some of the issues associated with replicating results across various fields.

Ryan (p. 1229) highlights the excitement and challenges that come with field-based research. In particular, observing processes as they occur in nature allows for discovery but makes replication difficult, because the precise conditions surrounding the observations are unique. Further, although laboratory research allows for the specification of experimental conditions, the conclusions may not apply to the real world. Debate about the merits of lab-based and field-based studies has been a persistent theme over time. Tomasello and Call (p. 1227) further contribute to this debate in their discussion of some obvious barriers to replication in primate cognition and behavior research (small numbers of subjects, expense, and ethics issues) as well as more subtle ones, such as the nontrivial challenge of designing tasks that elicit complex cognitive behaviors.

New technologies continue to produce a deluge of data of different varieties, raising expectations for new knowledge that will translate into meaningful therapeutics and insights into health. Ioannidis and Khoury (p. 1230) outline multiple steps for validating such large-scale data on the path to clinical utility and make suggestions for incentives (and penalties) that could enhance the availability of reliable data and analyses.

Peng (p. 1226) discusses the need for a minimum standard of reproducibility in computer sciences, arguing that enough information about methods and code should be available for independent researchers to reach consistent conclusions using original raw data. Specifically, he describes a model that one journal has used to make this a reality.

The need to convince the public that data are replicable has grown as science and public policy-making intersect, an issue that has has beset climate change studies. As Santer *et al.* (p. 1232) describe, having multiple groups examining the same data and generating new data has led to robust conclusions.

The importance of replication and reproducibility for scientists is unquestioned. Sometimes attempts to replicate reveal scientific uncertainties. This is one of the main ways that science progresses (see associated News stories of faster-than-light neutrinos and sirtuins, pp. 1200 and 1194). Unfortunately, in rare instances (compared to the body of scientific work), it can also indicate fraud (see the Editorial by Crocker and Cooper, p. 1182). How do we promote the publication of replicable data? The authors in this section come up with possibilities that are targeted at funders, journals, and the research culture itself. In the Readers' Poll, you can make your views known as well.

– BARBARA R. JASNY, GILBERT CHIN, LISA CHONG, SACHA VIGNIERI

Data Replication & Reproducibility

CONTENTS

Perspectives

See also Editorial p. 1182; News stories pp. 1194 and 1200; Readers' Poll p. 1203; Science Careers content p. 1179; and www.sciencemag.org/special/data-rep/

Science

Science

Replication
again, again, again

6 articles
9 pages

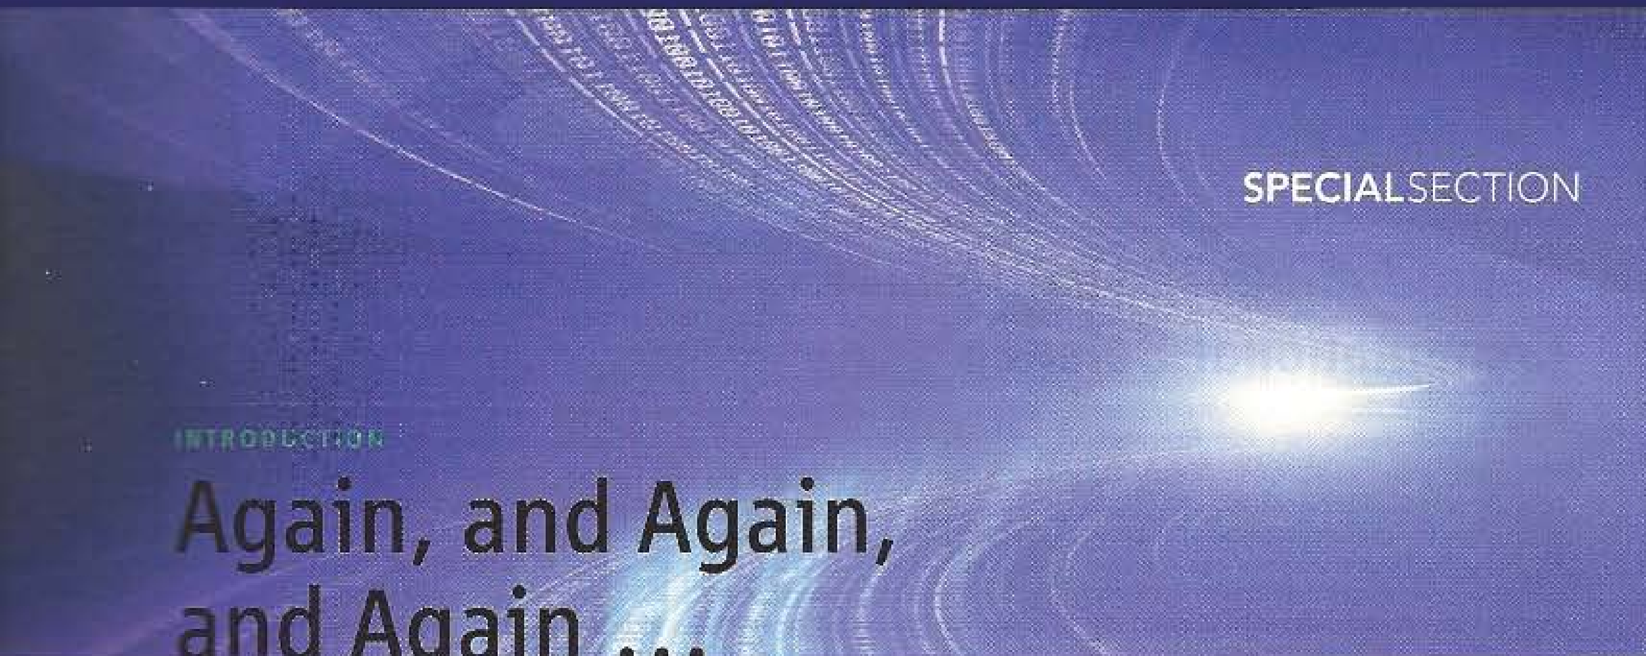

INTRODUCTION

Again, and Again, and Again …

REPLICATION—THE CONFIRMATION OF RESULTS AND CONCLUSIONS FROM ONE STUDY obtained independently in another—is considered the scientific gold standard. New tools and technologies, massive amounts of data, long-term studies, interdisciplinary approaches, and the complexity of the questions being asked are complicating replication efforts, as are increased pressures on scientists to advance their research. The five Perspectives in this section (and associated News and Careers stories, Readers' Poll, and Editorial) explore some of the issues associated with replicating results across various fields.

Ryan (p. 1229) highlights the excitement and challenges that come with field-based research. In particular, observing processes as they occur in nature allows for discovery but makes replication difficult, because the precise conditions surrounding the observations are unique. Further, although laboratory research allows for the specification of experimental conditions, the conclusions may not apply to the real world. Debate about the merits of lab-based and field-based studies has been a persistent theme over time. Tomasello and Call (p. 1227) further contribute to this debate in their discussion of some obvious barriers to replication in primate cognition and behavior research (small numbers of subjects, expense, and ethics issues) as well as more subtle ones, such as the nontrivial challenge of designing tasks that elicit complex cognitive behaviors.

New technologies continue to produce a deluge of data of different varieties, raising expectations for new knowledge that will translate into meaningful therapeutics and insights into health. Ioannidis and Khoury (p. 1230) outline multiple steps for validating such large-scale data on the path to clinical utility and make suggestions for incentives (and penalties) that could enhance the availability of reliable data and analyses.

Peng (p. 1226) discusses the need for a minimum standard of reproducibility in computer sciences, arguing that enough information about methods and code should be available for independent researchers to reach consistent conclusions using original raw data. Specifically, he describes a model that one journal has used to make this a reality.

The need to convince the public that data are replicable has grown as science and public policy-making intersect, an issue that has has beset climate change studies. As Santer *et al.* (p. 1232) describe, having multiple groups examining the same data and generating new data has led to robust conclusions.

The importance of replication and reproducibility for scientists is unquestioned. Sometimes attempts to replicate reveal scientific uncertainties. This is one of the main ways that science progresses (see associated News stories of faster-than-light neutrinos and sirtuins, pp. 1200 and 1194). Unfortunately, in rare instances (compared to the body of scientific work), it can also indicate fraud (see the Editorial by Crocker and Cooper, p. 1182). How do we promote the publication of replicable data? The authors in this section come up with possibilities that are targeted at funders, journals, and the research culture itself. In the Readers' Poll, you can make your views known as well.

– BARBARA R. JASNY, GILBERT CHIN, LISA CHONG, SACHA VIGNIERI

Data Replication & Reproducibility

CONTENTS

Perspectives

See also Editorial p. 1182; News stories pp. 1194 and 1200; Readers' Poll p. 1203; Science Careers content p. 1179; and www.sciencemag.org/special/data-rep/

Science

Science

Replication
again, again, again

6 articles
9 pages

It ~~was~~ is a core of statistics

INTRODUCTION

Again, and Again, and Again …

REPLICATION—THE CONFIRMATION OF RESULTS AND CONCLUSIONS FROM ONE STUDY obtained independently in another—is considered the scientific gold standard. New tools and technologies, massive amounts of data, long-term studies, interdisciplinary approaches, and the complexity of the questions being asked are complicating replication efforts, as are increased pressures on scientists to advance their research. The five Perspectives in this section (and associated News and Careers stories, Readers' Poll, and Editorial) explore some of the issues associated with replicating results across various fields.

Ryan (p. 1229) highlights the excitement and challenges that come with field-based research. In particular, observing processes as they occur in nature allows for discovery but makes replication difficult, because the precise conditions surrounding the observations are unique. Further, although laboratory research allows for the specification of experimental conditions, the conclusions may not apply to the real world. Debate about the merits of lab-based and field-based studies has been a persistent theme over time. Tomasello and Call (p. 1227) further contribute to this debate in their discussion of some obvious barriers to replication in primate cognition and behavior research (small numbers of subjects, expense, and ethics issues) as well as more subtle ones, such as the nontrivial challenge of designing tasks that elicit complex cognitive behaviors.

New technologies continue to produce a deluge of data of different varieties, raising expectations for new knowledge that will translate into meaningful therapeutics and insights into health. Ioannidis and Khoury (p. 1230) outline multiple steps for validating such large-scale data on the path to clinical utility and make suggestions for incentives (and penalties) that could enhance the availability of reliable data and analyses.

Peng (p. 1226) discusses the need for a minimum standard of reproducibility in computer sciences, arguing that enough information about methods and code should be available for independent researchers to reach consistent conclusions using original raw data. Specifically, he describes a model that one journal has used to make this a reality.

The need to convince the public that data are replicable has grown as science and public policy-making intersect, an issue that has has beset climate change studies. As Santer *et al.* (p. 1232) describe, having multiple groups examining the same data and generating new data has led to robust conclusions.

The importance of replication and reproducibility for scientists is unquestioned. Sometimes attempts to replicate reveal scientific uncertainties. This is one of the main ways that science progresses (see associated News stories of faster-than-light neutrinos and sirtuins, pp. 1200 and 1194). Unfortunately, in rare instances (compared to the body of scientific work), it can also indicate fraud (see the Editorial by Crocker and Cooper, p. 1182). How do we promote the publication of replicable data? The authors in this section come up with possibilities that are targeted at funders, journals, and the research culture itself. In the Readers' Poll, you can make your views known as well.

– BARBARA R. JASNY, GILBERT CHIN, LISA CHONG, SACHA VIGNIERI

Data Replication & Reproducibility

CONTENTS

Perspectives

See also Editorial p. 1182; News stories pp. 1194 and 1200; Readers' Poll p. 1203; Science Careers content p. 1179; and www.sciencemag.org/special/data-rep/

Science

Science again

Replication
again, again, again

6 articles
9 pages

4 Unfortunately, of statistics
it can also indicate fraud
-once upon a time !

INTRODUCTION

Again, and Again, and Again …

REPLICATION—THE CONFIRMATION OF RESULTS AND CONCLUSIONS FROM ONE STUDY obtained independently in another—is considered the scientific gold standard. New tools and technologies, massive amounts of data, long-term studies, interdisciplinary approaches, and the complexity of the questions being asked are complicating replication efforts, as are increased pressures on scientists to advance their research. The five Perspectives in this section (and associated News and Careers stories, Readers' Poll, and Editorial) explore some of the issues associated with replicating results across various fields.

Ryan (p. 1229) highlights the excitement and challenges that come with field-based research. In particular, observing processes as they occur in nature allows for discovery but makes replication difficult, because the precise conditions surrounding the observations are unique. Further, although laboratory research allows for the specification of experimental conditions, the conclusions may not apply to the real world. Debate about the merits of lab-based and field-based studies has been a persistent theme over time. Tomasello and Call (p. 1227) further contribute to this debate in their discussion of some obvious barriers to replication in primate cognition and behavior research (small numbers of subjects, expense, and ethics issues) as well as more subtle ones, such as the nontrivial challenge of designing tasks that elicit complex cognitive behaviors.

New technologies continue to produce a deluge of data of different varieties, raising expectations for new knowledge that will translate into meaningful therapeutics and insights into health. Ioannidis and Khoury (p. 1230) outline multiple steps for validating such large-scale data on the path to clinical utility and make suggestions for incentives (and penalties) that could enhance the availability of reliable data and analyses.

Peng (p. 1226) discusses the need for a minimum standard of reproducibility in computer sciences, arguing that enough information about methods and code should be available for independent researchers to reach consistent conclusions using original raw data. Specifically, he describes a model that one journal has used to make this a reality.

The need to convince the public that data are replicable has grown as science and public policy-making intersect, an issue that has has beset climate change studies. As Santer *et al.* (p. 1232) describe, having multiple groups examining the same data and generating new data has led to robust conclusions.

The importance of replication and reproducibility for scientists is unquestioned. Sometimes attempts to replicate reveal scientific uncertainties. This is one of the main ways that science progresses (see associated News stories of faster-than-light neutrinos and sirtuins, pp. 1200 and 1194). Unfortunately, in rare instances (compared to the body of scientific work), it can also indicate fraud (see the Editorial by Crocker and Cooper, p. 1182). How do we promote the publication of replicable data? The authors in this section come up with possibilities that are targeted at funders, journals, and the research culture itself. In the Readers' Poll, you can make your views known as well.

– BARBARA R. JASNY, GILBERT CHIN, LISA CHONG, SACHA VIGNIERI

Data Replication & Reproducibility

CONTENTS

Perspectives

See also Editorial p. 1182; News stories pp. 1194 and 1200; Readers' Poll p. 1203; Science Careers content p. 1179; and www.sciencemag.org/special/data-rep/

Science

Science again

Replication
again, again, again

6 articles
9 pages

Unfortunately, of statistics
it can also indicate fraud
- once upon a time !

Maybe !
" Fortunately
it can also indicate fraud "

SPECIALSECTION

INTRODUCTION

Again, and Again, and Again …

REPLICATION—THE CONFIRMATION OF RESULTS AND CONCLUSIONS FROM ONE STUDY obtained independently in another—is considered the scientific gold standard. New tools and technologies, massive amounts of data, long-term studies, interdisciplinary approaches, and the complexity of the questions being asked are complicating replication efforts, as are increased pressures on scientists to advance their research. The five Perspectives in this section (and associated News and Careers stories, Readers' Poll, and Editorial) explore some of the issues associated with replicating results across various fields.

Ryan (p. 1229) highlights the excitement and challenges that come with field-based research. In particular, observing processes as they occur in nature allows for discovery but makes replication difficult, because the precise conditions surrounding the observations are unique. Further, although laboratory research allows for the specification of experimental conditions, the conclusions may not apply to the real world. Debate about the merits of lab-based and field-based studies has been a persistent theme over time. Tomasello and Call (p. 1227) further contribute to this debate in their discussion of some obvious barriers to replication in primate cognition and behavior research (small numbers of subjects, expense, and ethics issues) as well as more subtle ones, such as the nontrivial challenge of designing tasks that elicit complex cognitive behaviors.

New technologies continue to produce a deluge of data of different varieties, raising expectations for new knowledge that will translate into meaningful therapeutics and insights into health. Ioannidis and Khoury (p. 1230) outline multiple steps for validating such large-scale data on the path to clinical utility and make suggestions for incentives (and penalties) that could enhance the availability of reliable data and analyses.

Peng (p. 1226) discusses the need for a minimum standard of reproducibility in computer sciences, arguing that enough information about methods and code should be available for independent researchers to reach consistent conclusions using original raw data. Specifically, he describes a model that one journal has used to make this a reality.

The need to convince the public that data are replicable has grown as science and public policy-making intersect, an issue that has has beset climate change studies. As Santer *et al.* (p. 1232) describe, having multiple groups examining the same data and generating new data has led to robust conclusions.

The importance of replication and reproducibility for scientists is unquestioned. Sometimes attempts to replicate reveal scientific uncertainties. This is one of the main ways that science progresses (see associated News stories of faster-than-light neutrinos and sirtuins, pp. 1200 and 1194). Unfortunately, in rare instances (compared to the body of scientific work), it can also indicate fraud (see the Editorial by Crocker and Cooper, p. 1182). How do we promote the publication of replicable data? The authors in this section come up with possibilities that are targeted at funders, journals, and the research culture itself. In the Readers' Poll, you can make your views known as well.

– BARBARA R. JASNY, GILBERT CHIN, LISA CHONG, SACHA VIGNIERI

Data Replication & Reproducibility

CONTENTS

Perspectives

See also Editorial p. 1182; News stories pp. 1194 and 1200; Readers' Poll p. 1203; Science Careers content p. 1179; and www.sciencemag.org/special/data-rep/

Science

Science again

Replication
again, again, again

6 articles
9 pages

Unfortunately, of statistics
it can also indicate fraud
once upon a time!

Maybe!
Fortunately
it can also indicate fraud

If fraud is there,
we ought to know about it!

6 Science 2013 Feb 15

Mr Bayes goes to Washington

Science 2013 Feb 15

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

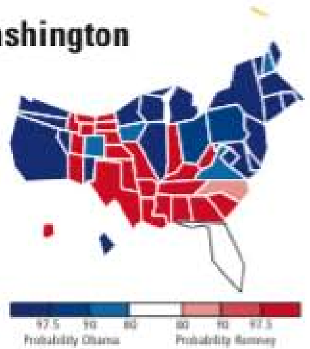

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Science 2013 Feb 15

Mr. Bayes Goes to Washington

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

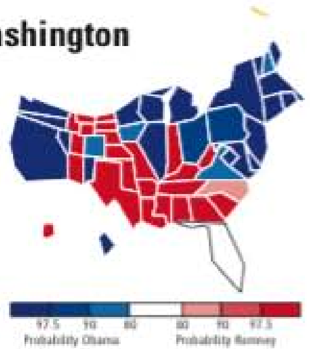

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

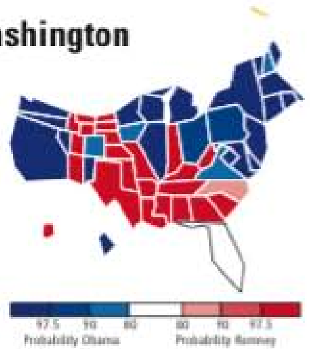

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. $27.95, C$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Science 2013 Feb 15

Mr. Bayes Goes to Washington

Glamor – Mr Bayes 1763
Washington DC

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

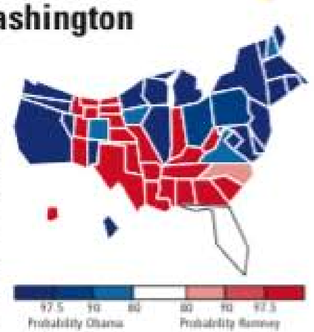

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Science 2013 Feb 15

Mr. Bayes Goes to Washington

Glamor – Mr Bayes 1763
Washington DC

Book review - by Wang & Campbell

Nate Silver: The Signal and The Noise

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

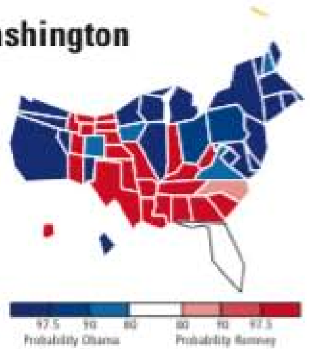

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

Science 2013 Feb 15

Mr. Bayes Goes to Washington

Glamor – Mr Bayes 1763
Washington DC

Book review – by Wang & Campbell

Nate Silver: The Signal and The Noise

- 2012 US Pres. election
- "aggregation of polls"

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (1). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (2)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (3). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (4). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (3–6) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

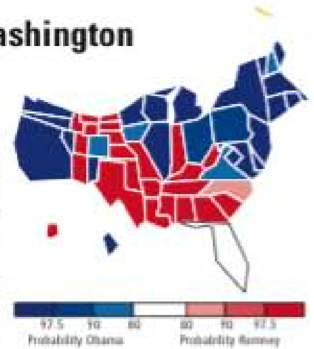

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. \$27.95, C\$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (7). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (8), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Science 2013 Feb 15

Mr. Bayes Goes to Washington

Glamor – Mr Bayes 1763
Washington DC

Book review – by Wang & Campbell

Nate Silver: The Signal and The Noise

- 2012 US Pres. election | highly successful!
- "aggregation of polls" | beat pundits!

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (1). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (2)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (3). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (4). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (3–6) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

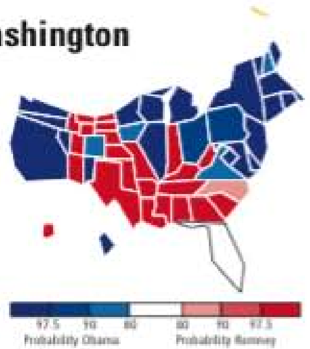

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. \$27.95, C\$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (7). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (8), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Science 2013 Feb 15

Mr. Bayes Goes to Washington

Glamor – Mr Bayes 1763
Washington DC

Book review – by Wang & Campbell

Nate Silver: The Signal and The Noise

- 2012 US Pres. election
- "Aggregation of polls"

| highly successful!
beat pundits!

but there's more ---

- Bayesian reasoning
- Modeling
- Fisher's prior
- fox & hedgehog

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

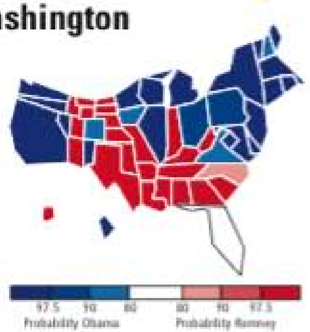

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. \$27.95, C\$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Science 2013 Feb 15

Mr. Bayes Goes to Washington

Glamor – Mr Bayes 1763
Washington DC

Book review – by Wang & Campbell

Nate Silver: The Signal and the Noise

- 2012 US Pres. election
- "Aggregation of polls"

highly successful!
beat pundits!

but there's more ---

- Bayesian reasoning (1)
- Modeling (2)
- Fisher's prior (3)
- fox & hedgehog (4)

Big issues for Statistics!

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

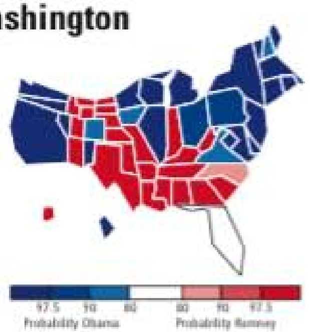

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

CREDIT: COURTESY SAM WANG/PRINCETON ELECTION CONSORTIUM

(1) Bayesian reasoning ? Bayes 1763

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

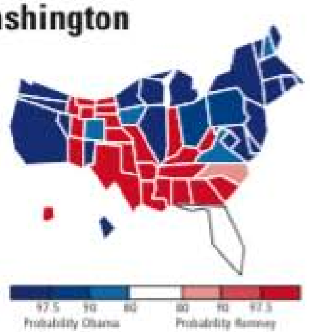

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

(1) Bayesian reasoning? Bayes 1763

- Bayes had Success/failure data on an unknown prob. p "Bin(n,p)"

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

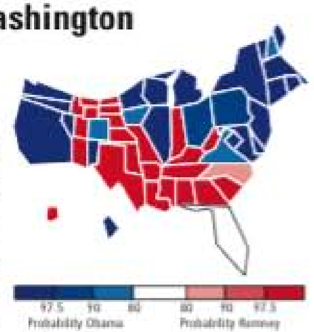

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. $27.95, C$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(1) Bayesian reasoning ? Bayes 1763

- Bayes had Success/failure data on an unknown prob. p "Bin(n,p)"

He imagined/supposed a roulette wheel or billiard table giving p Unif(0, 1)

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

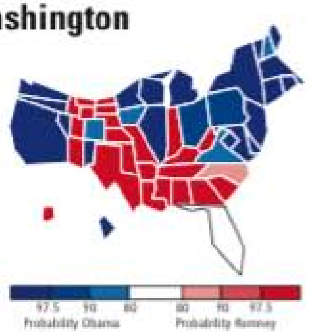

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. $27.95, C$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

(1) Bayesian reasoning? Bayes 1763

- Bayes had Success/failure data on an unknown prob. p "Bin(n,p)"
  He imagined/supposed a roulette wheel or billiard table giving p Unif$(0,1)$

- He combined them and used available "prob. theory from France"

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

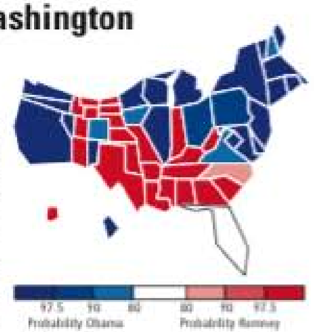

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(1) Bayesian reasoning ? Bayes 1763

- Bayes had Success/failure data on an unknown prob. p "$Bin(n,p)$"

He imagined/supposed a roulette wheel or billiard table giving p $Unif(0,1)$

- He combined them and used available "prob. theory from France"

Not quite what Silver promotes

- Bayes: missing data, so made it up

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

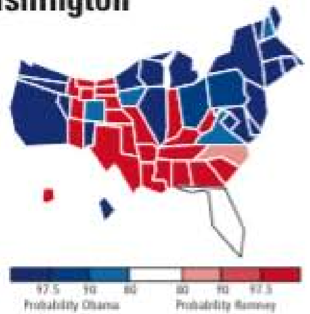

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

CREDIT: COURTESY SAM WANG/PRINCETON ELECTION CONSORTIUM

Published by AAAS

(1) Bayesian reasoning ? Bayes 1763

- Bayes had Success/failure data on an unknown prob. p "$Bin(n,p)$"

He imagined/supposed a roulette wheel or billiard table giving p $Unif(0,1)$

- He combined them and used available "prob. Theory from France"

Not quite what Silver promotes

- Bayes missing data, so made it up
  - Misrepresentation
  - Mysticism re Bayes

Efron - Science - 2013 June 7 genuine prior

(2) Modeling ?

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (1). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (2)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (3). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (4). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (3–6) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

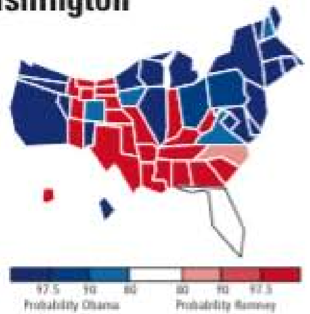

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (7). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (8), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

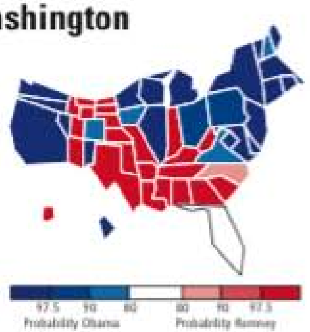

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

(2) Modeling?
Public - listen to the pundits
Silver - aggregate the polls

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (1). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (2)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (3). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (4). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (3–6) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

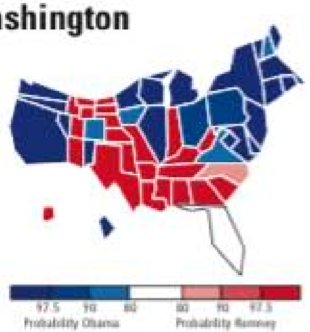

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. \$27.95, C\$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (7). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (8), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

(2) Modeling ?
- Public - listen to the pundits
- Silver - aggregate the polls
- Something quite sensible !

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

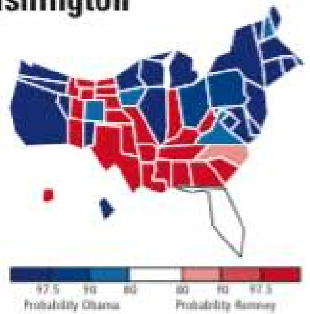

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(2) Modeling ?

- Public – listen to the pundits
- Silver – aggregate the polls
- Something quite sensible !

There is a whole subdiscipline of statistics on modeling !

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

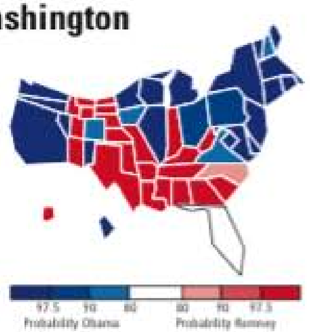

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(2) Modeling ?

- Public - listen to the pundits
- Silver - aggregate the polls
- Something quite sensible !

There is a whole subdiscipline of statistics on modeling !

Just standard "modeling"

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

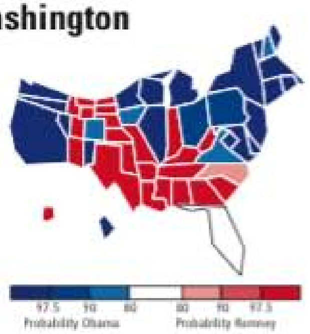

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(2) Modeling ?

- Public – listen to the pundits
- Silver – aggregate the polls
- Something quite sensible !

There is a whole subdiscipline of statistics on modeling !

Just standard "modeling"

Not Bayes

Not pundits

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

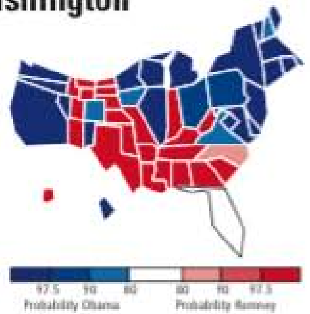

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(3) "Fisher's prior" concerning
<Smoking & lung cancer>

(3) "Fisher's prior" concerning

<Smoking & lung cancer>

Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the *New York Times* (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

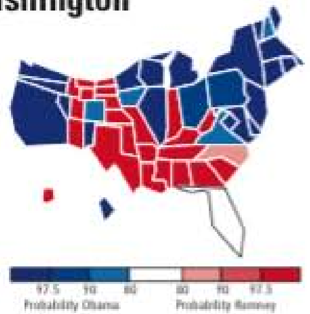

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/ The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

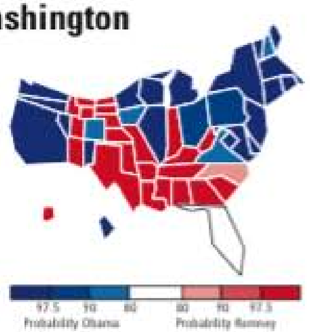

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. $27.95, C$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(3) "Fisher's prior" concerning
<Smoking & lung cancer>

Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

But Wang & Campbell fail to mention

(3) "Fisher's prior" concerning
<Smoking & lung cancer>

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

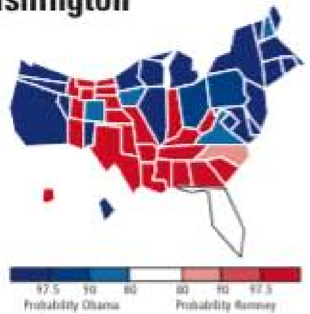

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012.
542 pp. $27.95, C$29.50.
ISBN 9781594204111.
Allen Lane, London. £25.
ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

But Wang & Campbell fail to mention
- Fisher - consultant - Imperial Tobacco

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111. Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

(3.) "Fisher's prior" concerning
< Smoking & lung cancer >

Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

But Wang & Campbell fail to mention
– Fisher – consultant – Imperial Tobacco
& found no experimental basis ...

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3–6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

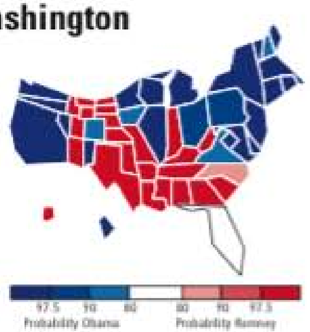

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't
The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

CREDIT: COURTESY SAM WANG/PRINCETON ELECTION CONSORTIUM

(3) "Fisher's prior" concerning
< Smoking & lung cancer >

Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

But Wang & Campbell fail to mention
- Fisher - consultant - Imperial Tobacco
& found no experimental basis ...
- Was: architect of Experimental Design

STATISTICS

Mr. Bayes Goes to Washington

Sam Wang[1] and Benjamin C. Campbell[2]

One day before the 2012 U.S. presidential election, former Reagan speechwriter Peggy Noonan wrote that "nobody knows anything" about who would win, asserting that Republican candidate Mitt Romney's supporters had the greater passion and enthusiasm (*1*). From a similarly data-free remove, columnist George Will predicted a Romney electoral landslide. MSNBC's Joe Scarborough said "it could go either way … anybody that thinks that this race is anything but a tossup right now … should be kept away from typewriters, computers, laptops, and microphones, because they're jokes." (*2*)

In the end, these pundits were the ones whose opinions proved dispensable. They were unable to detect a plain fact: based on public opinion polls with collectively excellent track records, President Obama had an advantage of 3 to 4 percentage points for nearly the entire campaign season. However, the world of political punditry measures success not by accuracy but by readership and viewership. And so it came to pass that legions of commentators expressed total confidence—and were wrong.

Beating the pundits has been possible since at least 2004, when one of us was among the first to statistically aggregate polls (*3*). In 2008, Nate Silver emerged as a poster child for aggregation, armed with a degree in economics, a love of numbers, and a professional track record in analyzing baseball performance and financial data. He enlivened a mostly suspenseless presidential race, providing timely quantitative analysis and color commentary on his website FiveThirtyEight, which became highly popular and was snapped up by the New York Times (*4*). His fame rose further in 2012, when he and other aggregators and modelers used hardnosed analysis (*3*–*6*) to silence skeptics.

Now Silver has written *The Signal and the Noise*, a book that addresses predictions not just in politics but in all aspects of modern life, with the eye of a hobbyist and a sense of fun. Freed from the word limits of blog essays, the book is a meandering, nerd's-eye view of what principles, if any, are common to good forecasting in daily life, leisure activity, and science.

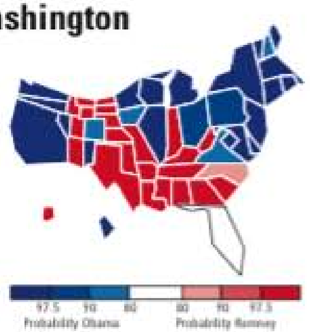

Validated by the outcome. The Princeton Election Consortium's final electoral college predictions for November 2012. (States are sized according to their share of electoral votes.)

We use predictions to guide our future actions, from planning weekend outings to taking care of our health, but most people have no idea how scientific predictions are made. This book is for them. Silver introduces some of the concepts behind data modeling, including probability, Bayesian inference, and uncertainty. He takes lengthy looks at topics ranging from flu epidemics to the 1996 chess-playing triumphs of Deep Blue.

The Signal and the Noise
Why So Many Predictions Fail—But Some Don't/The Art and Science of Prediction
by Nate Silver
Penguin, New York, 2012. 542 pp. $27.95, C$29.50. ISBN 9781594204111.
Allen Lane, London. £25. ISBN 9781846147524.

A reappearing theme in *The Signal and the Noise* is Bayesian reasoning, an approach that has swept the sciences. Probability had been conventionally interpreted as meaning the true likelihood of an event—for instance, how often the total of two rolled dice will add up to seven. Such a "frequentist" point of view has in many cases given way to an approach pioneered by Reverend Thomas Bayes in the 18th century, which emphasizes that probability can only be interpreted in terms of the hypotheses that preceded the measurement.

Although Silver asserts that Bayesian political forecasting has more in common with poker than with hard sciences such as physics and biology, these topics all use the same mathematical toolkit. Large-scale physics collaborations depend on sensitive models to predict the probabilistic decay rates of particles, looking for outliers that might represent signals in the noise and hence discoveries. In our field, many neuroscientists have begun to view the brain as a prediction machine (*7*). We perceive the world around us by making inferences from noisy and incomplete data. To do so, the brain must form a model of its environment—a set of "priors" learned over a lifetime that is used to interpret incoming data. This Bayesian machine continually updates its priors to correspond to its environment. Through this process, our brains spend many years honing appropriate priors for the complex tasks that we perform effortlessly.

Silver gives a well-known equation for how to take into account the Bayesian prior but doesn't show where it comes from. Readers wanting a deeper explanation of Bayes's rule might consult another source such as BetterExplained.com (*8*), which teaches the subject by using e-mail spam filtering as an example. Silver's chosen anecdotes include the classic example of mammogram interpretation—but also how to interpret that unfamiliar underwear that just showed up in your partner's dresser drawer.

At times Silver writes as if the cure for bad modeling can be reduced to "more Bayes." Such a prescription does not do justice to the historic controversies surrounding interpretations of probability. A beginner might come away from this book believing that an earlier generation of frequentists were simply ignorant. In a cartoonish account, Silver lobs a broadside at a monumental figure in statistics, Ronald A. Fisher, who late in life argued against the idea that smoking causes cancer—and who coined "Bayesian" as a derogatory term. Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (*9*). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

Our biggest criticism of the book is that although statistics and Bayesian inference are powerful ideas, they are not a cure-all. In his enthusiasm for the good Reverend, Silver has stuffed a fair bit into the same Procrustean bed. Silver uses the old fox-hedgehog analogy, saying that foxes (including himself) use many ideas, whereas hedgehogs focus on one subject only. But here he is a hedgehog with one big idea: statistics.

However, Bayesian reasoning works only if the prior is adapted for the task. According to Silver, many of today's "half-baked

[1]Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA. E-mail: sswang@princeton.edu [2]Laboratory of Biological Modeling, Rockefeller University, 1230 York Avenue, New York, NY 10065, USA. E-mail: bcampbell@rockefeller.edu

Published by AAAS

(3) "Fisher's prior" concerning
< Smoking & lung cancer >

Silver suggests that Fisher's aversion to Bayes caused him to err. In fact, the real problem was that Fisher was a smoker (9). Fisher's prior beliefs prevented him from accepting epidemiological and biological evidence, an erroneous prior if ever there was one.

But Wang & Campbell fail to mention
- Fisher - consultant - Imperial Tobacco
& found no experimental basis ...
- Was: architect of Experimental Design
- which verified Higgs boson

Now - Climate change
- Evidence based medicine
- Vioxx

Not so simple !

7 Science 2013 June 7

Efron: Bayes Theorem

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6*, *9*, *10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5*, *7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al* ... duction of su ... taxonomically ... teria from a r ... Some of their ... isolates; these ... erate superoxi ... the vast expa ... not receive su ... bacteria are n ... and may thus ... ciated biologi ... to the ocean surface (see the figure).

"controversial theorem"

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8*, *9*, *13*, *14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and ... ity, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).

ONE **5**, e12508 (2010).

12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory. [redacted]

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but
- just "conditional probability"

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
- just "conditional probability"
- been around for a 100 years

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
- just "conditional probability"
- been around for a 100 years
- Just the application was different

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
- just "conditional probability"
- been around for a 100 years
- Just the application was different
- Bayes: imagined a prior experiment

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
  - just "conditional probability"
  - been around for a 100 years
  - Just the application was different
- Bayes: imagined a prior experiment
- Laplace: proposed a mathematical input

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
  - just "conditional probability"
  - been around for a 100 years
  - Just the application was different
- Bayes: imagined a prior experiment
- Laplace: proposed a mathematical input
- Savage: How he felt, would bet

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
  - just "conditional probability"
  - been around for a 100 years
  - Just the application was different
- Bayes: imagined a prior experiment
- Laplace: proposed a mathematical input
- Savage: How he felt, would bet
- Jeffreys: For convenience

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
  - just "conditional probability"
  - been around for a 100 years
  - Just the application was different
- Bayes: imagined a prior experiment
- Laplace: proposed a mathematical input
- Savage: How he felt, would bet
- Jeffreys: For convenience

Not Theorem

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6*, *9*, *10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5*, *7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8*, *9*, *13*, *14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
  - just "conditional probability"
  - been around for a 100 years
  - Just the application was different
- Bayes: imagined a prior experiment
- Laplace: proposed a mathematical input
- Savage: How he felt, would bet
- Jeffreys: For convenience

Not ~~Theorem~~ What you do with "condn'l prob."

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6, 9, 10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5, 7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8, 9, 13, 14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Bayes' 1763 paper was an impeccable exercise in probability theory.

- but ...
  - just "conditional probability"
  - been around for a 100 years
  - Just the application was different
- Bayes: imagined a prior experiment
- Laplace: proposed a mathematical input
- Savage: How he felt, would bet
- Jeffreys: For convenience

Not ~~Theorem~~ What you do with "cond'l prob."

Mis-use of language!

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (6, 9, 10). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (5, 7). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (8, 9, 13, 14). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (12). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (14). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (9).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, M. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (1). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (2).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (2), one having only neutral effects on the output of Bayes' rule (3). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (6, 9, 10). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (5, 7). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (8, 9, 13, 14). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (12). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (14). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (9).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (1). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (2).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (2), one having only neutral effects on the output of Bayes' rule (3). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

[redacted] in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (6, 9, 10). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (5, 7). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (8, 9, 13, 14). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (12). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (14). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (9).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, M. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (1). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (2).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (2), one having only neutral effects on the output of Bayes' rule (3). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

Laplace 1749-1827 pre-eminent mathematician

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6*, *9*, *10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5*, *7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8*, *9*, *13*, *14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240199

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

Laplace 1749-1827 pre-eminent mathematician

Another view:

Laplace pre-viewed confidence (Fisher 1930)

(Stat. Science 2011 p 299-316)

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (6, 9, 10). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (5, 7). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (8, 9, 13, 14). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (12). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (14). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (9).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (1). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (2).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (2), one having only neutral effects on the output of Bayes' rule (3). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

Laplace 1749-1827 pre-eminent mathematician

Another view:

Laplace pre-viewed confidence (Fisher 1930)

(Stat. Science 2011 p 299-316)

Brad then, in "own practice",

.. uses Bayes only in presence of

genuine prior information.

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (6, 9, 10). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (5, 7). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (8, 9, 13, 14). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (12). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (14). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (9).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (1). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (2).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (2), one having only neutral effects on the output of Bayes' rule (3). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

Laplace 1749-1827 pre-eminent mathematician

Another view:

Laplace pre-viewed confidence (Fisher 1930)

(Stat. Science 2011 p 299-316)

Brad then, in "own practice",

..uses Bayes only in presence of

genuine prior information.

..but such is just "modeling" not "Bayes"

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6*, *9*, *10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5*, *7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8*, *9*, *13*, *14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

Laplace 1749-1827 pre-eminent mathematician

Another view:

Laplace pre-viewed confidence (Fisher 1930)

(Stat. Science 2011 p 299-316)

Brad then, in "own practice",

..uses Bayes only in presence of

genuine prior information.

..but such is just "modeling" not "Bayes"

Many Bayesians use Laplace priors

resentatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (*6*, *9*, *10*). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (*5*, *7*). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (*8*, *9*, *13*, *14*). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (*12*). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (*14*). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (*9*).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafiriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. I* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

Acknowledgments: Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

MATHEMATICS

Bayes' Theorem in the 21st Century

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

Bradley Efron

The term "controversial theorem" sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (*1*). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys. They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on a daily basis, correctly predicting the actual vote in all 50 states. "Statisticians beat pundits" was the verdict in the press (*2*).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an "uninformative prior" or "objective prior" (*2*), one having only neutral effects on the output of Bayes' rule (*3*). Whether or not this

Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

Brad: "No trouble in presence of genuine prior info"

But...

in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator.

Laplace 1749-1827 pre-eminent mathematician

Another view:

Laplace pre-viewed confidence (Fisher 1930)

(Stat. Science 2011 p 299-316)

Brad then, in "own practice",

..uses Bayes only in presence of

genuine prior information.

..but such is just "modeling" not "Bayes"

Many Bayesians use Laplace priors

but "escape" Brad .... or try to ...

by calling them objective (...) priors

Directions? We've seen: "There are serious risks!"

What is being done?

In statistics there are multiple theories! What gives?

Directions? We've seen: "There are serious risks!"
What is being done?

In statistics there are multiple theories! What gives?

I'll mention:

1) All <u>relevant</u> information

Directions? We've seen: "There are serious risks!"

What is being done?

In statistics there are multiple theories! What gives?

I'll mention:

1) All <u>relevant</u> information

2) Where theta is the data: <u>Higher order</u>

Directions? We've seen: "There are serious risks!"

What is being done?

In statistics there are multiple theories! What gives?

I'll mention:

1) All relevant information

2) Where theta is re data: Higher order

3) Bootstrap ≡ Higher order

Directions? We've seen: "There are serious risks!"

What is being done?

In statistics there are multiple theories! What gives?

I'll mention:

1) All <u>relevant</u> information

2) Where theta is re data: <u>Higher order</u>

3) <u>Bootstrap</u> ≡ Higher order

4) Good <u>Bayes</u> ⇒ Approx <u>Confidence</u>

1) All relevant information ?

1) All relevant information ?

Statistical tradition : use a sufficient statistic !

1) All relevant information ?

Statistical tradition : use a sufficient statistic !

Trouble: they rarely exist ! ... Normality

1) All relevant information ?

Statistical tradition: use a sufficient statistic!

Trouble: they rarely exist! ... Normality

But recent theory (regular model) has developed:

- With Interest parameter $\psi(\theta)$
- Essentially unique dist'n for assessing ψ

1) All relevant information ?

Statistical tradition : use a sufficient statistic !

Trouble: they rarely exist ! ... Normality

But recent theory (regular model) has developed :

- With Interest parameter $\psi(\theta)$
- Essentially unique dist'n for assessing ψ

$$\frac{e^{k/n}}{(2\pi)^{d/2}} \exp\left\{-\frac{r^2(\psi; s)}{2}\right\} \left|\hat{j}_{\varphi\varphi}(s)\right|^{-1/2} \left|\hat{j}_{(\lambda\lambda)}(s)\right|^{1/2} .$$

1) All <u>relevant</u> information?

Statistical tradition: use a <u>sufficient</u> statistic!

Trouble: they <u>rarely exist</u>! ... Normality

But recent theory (regular model) has developed:

- With Interest parameter $\psi(\theta)$
- Essentially <u>unique</u> dist'n for assessing ψ

$$\frac{e^{k/n}}{(2\pi)^{d/2}} \exp\left\{-\frac{r^2(\psi;s)}{2}\right\} |\hat{j}_{\varphi\varphi}(s)|^{-1/2} |\hat{j}_{(\lambda\lambda)}(s)|^{1/2} .$$

- Simple ingredients; immediate
- Here? Just to indicate it is <u>real</u>!

2) Where theta is the data: Higher order

2) Where theta is ne data: Higher order

a whole area of statistics ... since 1954

-Uses model continuity !

2) Where theta is the data: Higher order

a whole area of statistics ... since 1954 (SP)

- Uses model continuity!

- Uses all relevant information (preceding)

2) Where *theta* is re *data*: Higher order

A whole area of statistics ... since 1954 (SP)

- Uses model continuity !
- Uses all relevant information (preceding)

Gives precise statistical position: "data re-theta"

- $p = \Phi(r^*)$
- Availability !

3) Bootstrap ≡ Higher order

3) Bootstrap ≡ Higher order

Statistics has multiple theories:

frequentist; Bayes; Higher order; Bootstrap; ..

3) <u>Bootstrap</u> $\equiv$ Higher order

Statistics has multiple theories:

frequentist; Bayes; Higher order; Bootstrap; ..

How can statistics have conflicting theories?

3) Bootstrap $\equiv$ Higher order

Statistics has multiple theories:

frequentist; Bayes; Higher order; Bootstrap; ..

How can statistics have conflicting theories?

BS: Sample from $f(y;\hat{\theta}^0)$

3) Bootstrap $\equiv$ Higher order

Statistics has multiple theories:

frequentist; Bayes; Higher order; Bootstrap; ..

How can statistics have conflicting theories?

BS: Sample from $f(y; \hat{\theta}^0)$

Modify: " " $f(y; \hat{\theta}^0_\psi)$

3) <u>Bootstrap</u> $\equiv$ Higher order

Statistics has multiple theories:

frequentist; Bayes; Higher order; Bootstrap; ..

How can statistics have conflicting theories?

BS: Sample from $f(y; \hat{\theta}^0)$

Modify: " " $f(y; \hat{\theta}^0_\psi)$

$\Rightarrow$ Bootstrap $\equiv$ Higher order! New!

3) Bootstrap ≡ Higher order

Statistics has multiple theories:

frequentist; Bayes; Higher order; Bootstrap; ..

How can statistics have conflicting theories?

BS: Sample from $f(y;\hat{\theta}^0)$

Modify: " " $f(y;\hat{\theta}^0_\psi)$

then Bootstrap same as Higher order! New!

but ·09 sec vs. 20 hours

4 Good Bayes ⇒ Approx Confidence

4 Good Bayes $\Rightarrow$ Approx Confidence

Bayes is heavily promoted by pro-Bayesian

4 Good Bayes ⇒ Approx Confidence

Bayes is heavily promoted by pro-Bayesians

- Use all relevant information ... see ① above "L"

4. Good Bayes ⇒ Approx Confidence

Bayes is heavily promoted by pro-Bayesians:

- Use all relevant information --- see (1) above

- Use a prior

4 Good Bayes ⇒ Approx Confidence

Bayes is heavily promoted by pro-Bayesians:

- Use all relevant information --- see (1) above

- Use a prior a) genuine prior info ... Should have been in model

4 Good Bayes $\Rightarrow$ Approx Confidence

Bayes is heavily promoted by pro-Bayesians:

- Use all relevant information ... see (1) above

- Use a prior a) genuine prior info ... Should have been in model

  b) Un-informative .. Gives approximate confidence

4 Good Bayes ⇒ Approx Confidence

Bayes is heavily promoted by pro-Bayesians:

- Use all relevant information ... see ① above

- Use a prior a) genuine prior info ... Should have been in model

  b) Un-informative .. Gives approximate confidence

  c) Betting judgment ... Report separately for user to be cautious

Conclusions

Conclusions

Risks for statistics;

Responsibility also for statistics

Conclusions

Risks for statistics

Responsibilities... for statistics

Multiple theories: How can a discipline tolerate them?

Physics seeks $b's from taxpayers to resolve such

Conclusions

Risks for statistics

Responsibilities... for statistics

Multiple theories: How can a discipline tolerate them?

Physics seeks $b's from taxpayers to resolve such

But: Bootstrap & frequentist (Higher order) are equivalent

Conclusions

Risks for statistics

Responsibilities... for statistics

Multiple theories: How can a discipline tolerate them?

Physics seeks \$b's from taxpayers to resolve such

But: Bootstrap & frequentist (Higher order) are equivalent

Bayes? - Use genuine prior information --- Efron 2013

- then it is just frequentist "

Conclusions

Risks for statistics

Responsibilities... for statistics

Multiple theories: How can a discipline tolerate them?

Physics seeks \$b's from taxpayers to resolve such

But: Bootstrap & frequentist (Higher order) are equivalent

Bayes? - Use genuine prior information --- Efron 2013

- then it is just frequentist "

- Laplace type un-informative (developed version)

- Just approximate confidence

Conclusions

Risks for statistics

Responsibilities... for statistics

Multiple theories: How can a discipline tolerate them?

Physics seeks $b's from taxpayers to resolve such

But: Bootstrap & frequentist (Higher order) are equivalent

Bayes? - Use <u>genuine</u> prior information --- Efron 2013

- then it is just frequentist "

- Laplace type un-informative (developed version)

- Just <u>approximate confidence</u>

Statistics is <u>strong</u>

- but has deep <u>responsibilities</u>

1 Data can overwhelm statistics ?

2 Clinical trials "can" be replaced by more Data ?

3. Vioxx | Statistics overlooked

4. L'aquila |

Thank you!

5. Replication Needed but neglected

6. Bayes | Call conditional prob. by another name

7. Bayes | and create a lot of mystery

Directions : All relevant information

Where theta is re data : Higher order

Bootstrap ≡ Higher order ·09 sec vs. 20 hours

Good Bayes ⇒ Approx Confidence