

proved why back-to-back titles have been so hard to achieve, falling into an epic collapse: he bogeyed holes 10 and 11, and then *quadruple* bogeyed the 12th by sending two balls into the water. In the end it was Danny Willett, a relatively unheralded Englishman, who held on to claim his first major championship at five under par.

Upon witnessing such a remarkable implosion, *The Economist's* history- and statistics-minded golf fans immediately wondered where it ranked among the sport's all-time chokes. Much to our surprise, we were not able to find a readily available online database of golfers' mathematical probability of victory at every point in a tournament. (If any well-informed readers happen to be aware of one that we missed, we will gladly link to it.) Although Ken Pomeroy, a former meteorologist better-known for his college-basketball rankings, has built a robust golf-forecasting model, its past predictions reside only on a Twitter feed. Mr Pomeroy himself has written that "if you are really interested in developing something better, it's totally possible". With extreme hubris, we have decided to take up his challenge.

We are pleased to announce the release of EAGLE, the *Economist* Advantage in Golf Likelihood Estimator (backronyms have become de rigueur in the sports-forecasting world). EAGLE is a mathematical model of golf tournaments that estimates every player's chances of victory at every point in the event. A detailed description of how it works sits at the end of this post.

Advertisement



Latest updates »



Diversification in Africa: Gabon's struggle to move beyond oil

Middle East and Africa | 1 hour 39 mins ago



You've got a nerve: Researchers have built an artificial neuron

Science and technology | 1 hour 57 mins ago



The Economist explains: How Clinton and Trump plan to boost wages

The Economist explains | Aug 3rd, 04:45



Daily chart: Rio's long decline

Graphic detail | Aug 2nd, 17:04



Money talks: Stressed out banks

Free exchange | Aug 2nd, 16:42



Germany's Turkish minority: Tensions in Turkey are spilling over to...

Europe | Aug 2nd, 16:40



The Petrobras scandal: A judge indicts Brazil's former president

The Americas | Aug 2nd, 16:01

More latest updates »

Most commented



1 Globalisation and politics
The new political divide

- 2 Putin, Trump and the DNC:** Signal and noise
- 3 Hillary Clinton on stage:** The Democrats' orchestral finale
- 4 American foreign policy:** Playing it long
- 5 How Germans handle terror:** Pure reason

Advertisement

A young millennium of golf, hole by hole

Selected win probabilities in major men's tournaments

Choose year and tournament to see relevant results. Hover over key to isolate player

2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016

MASTERS US OPEN BRITISH OPEN PGA CHAMPIONSHIP

Predicted winner at start

Jason Day

Max. win probability (finishing 3rd or lower)

Jason Day

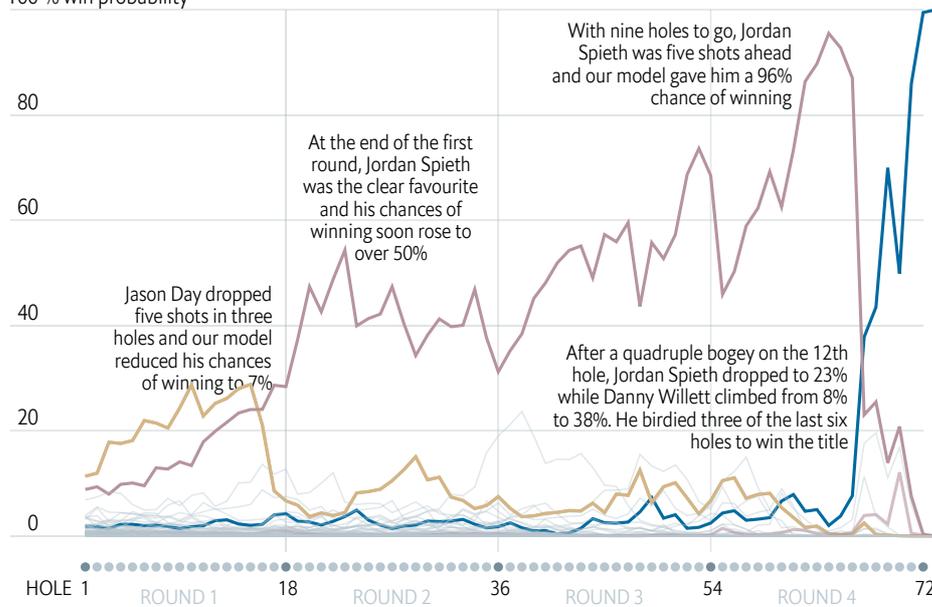
Actual winner

Danny Willett

Actual second place

Jordan Spieth
Lee Westwood

100 % win probability



Source: *The Economist*

But for fans who simply want to get to the punch line, we are publishing the system's results here. Above are win-probability graphs of every major men's tournament since 2001 that EAGLE was able to forecast (we hope to fill in the handful of missing ones soon). You can see how the likelihood of victory evolved over the course of the event for the champion, all runners-up, the golfer we would have projected to win before play began and the golfer who attained the highest chance of victory before finishing third or lower.

The numbers show that not only was Mr Spieth's collapse not "the most shocking in golf history", as one ESPN writer claimed, it wasn't even the most shocking this decade. In the 2012 British Open, Adam Scott had a four-stroke lead with just four holes left to play—good for a 98.6% win probability—and somehow managed to implode. But Mr Spieth's 2016 Masters performance was indeed the second-most unlikely in our dataset, as EAGLE gave him a 95.5% chance of victory (see chart 1). The only other player to reach above 90% and lose was Kenny Perry, who squandered a two-stroke edge with two holes left during the 2009 Masters.

For readers following the 2016 Open Championship, which runs from July 14th to 17th, EAGLE's projected leaderboard sits at the top of this page. It lists the ten players we believe should be favoured, based on EAGLE's estimates of their skill, the difficulty of the Royal Troon course and the current standings. We will update these forecasts at regular intervals throughout the event. Before the first tee, EAGLE saw Jason Day, the current world number one, as the front-runner, followed by the 2016 US Open winner Dustin Johnson, Mr Spieth and Rory McIlroy. But tournaments this large (156 players will tee off at The Open) are extraordinarily hard to predict, and the most likely outcome was that

Whatever the χ discover the γ

Enjoy 12 weeks access for €57

LEARN MORE

Economist blogs

- Bagehot's notebook | British politics
- Buttonwood's notebook | Financial markets
- Democracy in America | American politics
- Erasmus | Religion and public policy
- Free exchange | Economics
- Game theory | Sports
- Graphic detail | Charts, maps and infographics
- Gulliver | Business travel
- Prospero | Books, arts and culture
- The Economist explains | Explaining the world, daily

Products and events

Test your EQ

Take our weekly news quiz to stay on top of the headlines

Want more from *The Economist*?

Visit The Economist e-store and you'll find a range of carefully selected products for business and pleasure, Economist books and diaries, and much more

Advertisement

Whatever the χ discover the γ

Enjoy 12 weeks access for €57

LEARN MORE

Look out below

1

Worst collapses in major men's golf tournaments, 2001-16

Golfer	Tournament	Max win prob probability	Max win prob at:		Lead in strokes at max win prob	Finish	Champion
			round	hole			
Adam Scott	2012 British Open	0.986	4	15	4	2nd	Ernie Els
Jordan Spieth	2016 Masters	0.955	4	10	5	2nd	Danny Willett
Kenny Perry	2009 Masters	0.910	4	17	2	2nd	Ángel Cabrera
Jason Dufner	2011 PGA	0.869	4	16	2	2nd	Keegan Bradley
Len Mattiace	2003 Masters	0.806	4	18	1	2nd	Mike Weir
Sergio García	2007 British Open	0.779	4	4	4	2nd	Padraig Harrington
Tiger Woods	2009 PGA	0.772	4	12	1	2nd	Y.E. Yang
Dustin Johnson	2010 PGA	0.761	4	18	1	5th	Martin Kaymer
Ernie Els	2004 Masters	0.755	4	16	1	2nd	Phil Mickelson
Tom Watson	2009 British Open	0.720	4	18	1	2nd	Stewart Cink

Source: *The Economist*

Economist.com

none of those four favourites will win. Their combined probability of victory was just 28% according to EAGLE, and 33% according to the punters at [Betfair Exchange](#) (BE).

You may be wondering whether you can use EAGLE to make money. We wouldn't recommend trying. Using historical betting lines from [GolfOdds.com](#), we compared EAGLE's past pre-tournament projections to those of leading bookmakers. Although the two generally tracked quite closely—their [correlation coefficient](#) was 0.84, where 1 means moving in perfect unison—we found both that the bookies' lines were more accurate than ours, and that adding in a dose of EAGLE to the market predictions did not improve their forecasts. We're not discouraged by this: given that EAGLE relies exclusively on publicly available course, scorecard and ranking data, it would be quite striking if highly competitive betting markets had not already incorporated such information fully. What EAGLE can do that Las Vegas can't is tell you not just what will happen, but why. And it's still possible that EAGLE's ability to update its forecasts as an event proceeds might outpace the markets. We plan to collect data on this during future tournaments.

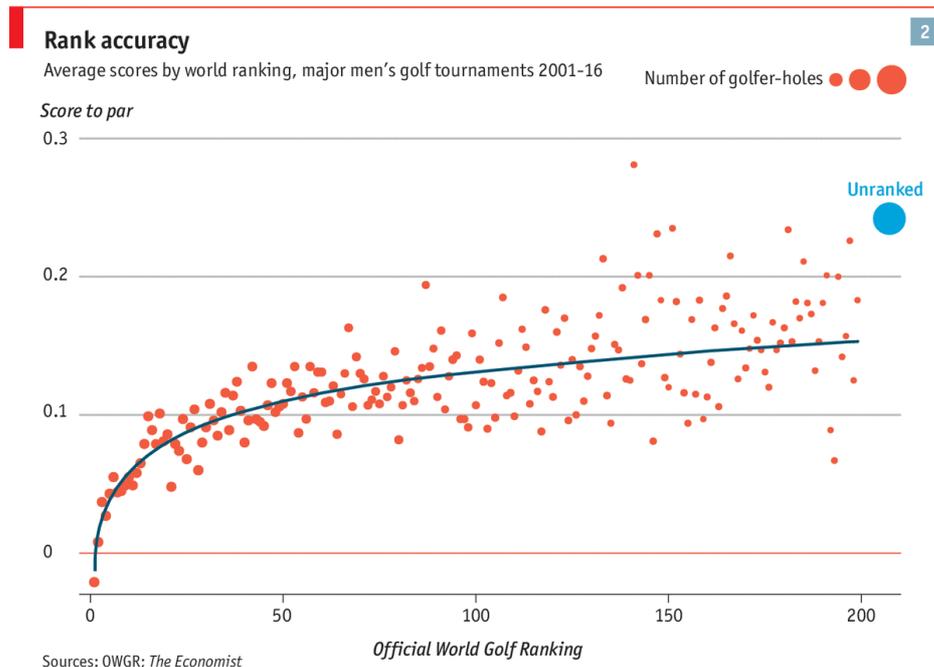
Nonetheless, despite the resounding evidence that EAGLE in its current state does not beat the bookies, we're still inclined to throw some money away as a demonstration of our faith in the model. We've placed wagers on three players the system is particularly fond of relative to the betting markets. Before The Open began, EAGLE thought [Bubba Watson](#) was about a 47-to-one shot. But he was available on BE at 84-to-one before fees. EAGLE put [Jim Furyk](#) at 60-to-one, while BE was selling him at 100-to-one. And the ultra-longshot [Jeung-hun Wang](#) got 233-to-one from EAGLE, and a whopping 1,000-to-one at BE. We're splitting a £100 bet among these contenders. By our own estimation, there is a 95.8% chance we will lose. Wish us luck.

Methodology: Fly like an eagle

EAGLE is based on a dataset of some 440,000 golfer-holes drawn from major tournaments during the past 15 years. Conceptually, it is quite straightforward. First, the system estimates each player's skill level on every hole they play. Next, it calculates the probability of a golfer of that ability producing each score type on every hole, ranging from an eagle (two strokes below par) to a quadruple bogey (four strokes above it), according to the hole's configuration and difficulty. Finally, it uses those probabilities to simulate every hole left to play in the event 10,000 times for each golfer, counting the number of simulations won by each participant in order to derive their chances of victory.

EAGLE's number-crunching starts with the simplest of premises: that the [Official World Golf Ranking](#) (OWGR) is a reasonable proxy for a player's skill. The evidence supports this hypothesis. Players with worse ranks tend to have inferior scores to those with better ranks. Those who do not make the top 200 at all fare even poorer (see chart 2). However, the OWGR is a far more reliable guide to the ability of elite players, whose scores are

clustered around par, than it is to that of golfers in the bottom half of the rankings, whose average results sometimes push much closer towards bogey range and sometimes don't. Because of this phenomenon, we have nicknamed this graph the "gramophone plot".



Economist.com

The next step in the analysis also stands to reason: that holes of different par values are likely to display different characteristics, and that holes where the pin sits far from the tee will be more difficult than those where it is close. The numbers bear out this assumption as well. Although absolute stroke totals are of course highest on par-fives and lowest on par-threes, par tends to be much easier to beat on par-fives than on fours and threes (see chart 3). At the same time, scores on par fives are also much more influenced by the distance to the hole than those on fours and threes are. Presumably, this is because only the top drivers in the sport can reach a far-off green fast enough to have a shot at beating par. In addition, scores on par-threes tend to cluster very close to par with a smattering of bogeys, whereas birdies and even eagles are far more common on par-fives. So a player seeking to protect a narrow lead will benefit from a course with a lot of par-threes, which prevent opponents from making up too much ground. One hoping for an unlikely comeback should hope for a steady diet of par-fives.

Together, these three factors—the OWGR, par and distance—account for a sizeable chunk of the variation in golfers' scores. By combining them in a multiple linear regression, EAGLE can calculate a baseline forecast score for any golfer on any hole. For example, we'd expect a player ranked fifth in the world on a relatively short 500-yard (457-metre) par-five to score around 0.407 strokes below par—say, 59.3% pars and 40.7% birdies. In contrast, an unranked player on a challenging par-four of the same length would be projected for a score of 0.479 strokes *above* par, perhaps in the form of 52.1% pars and 47.9% bogeys.

Everyone's a special snowflake

Unfortunately, this approach relies on two false, simplifying assumptions: that all holes of the same par and distance are alike, and that all golfers of the same world ranking are alike. Let's dispel with these fictions.

Although we don't have specific historical data about non-distance factors that contribute to a hole's difficulty, such as bunkers and water hazards, we can infer information about

All of golf is divided into three parts

Average scores and distance, major men's golf tournaments 2001-16

Number of golfer-holes ● ● ●



Source: *The Economist*

Economist.com

them from how golfers perform as a group. The first time we see a course, we assume that each hole displays normal difficulty for its par and distance. By the second round, we can compare the average score of everyone who played it previously with the result we would have expected based on their OWGRs, and on the holes' pars and yardages. On some holes, seemingly everyone in the field will birdie; on others even a straight par looks like a remarkable achievement.

These patterns turn out to be highly consistent within tournaments, though they are not particularly durable from year to year. Take the notorious [14th hole at Pebble Beach](#) in California (pictured), a 580-yard par-five, which was last played in a major for the 2010 US Open. Based on the OWGRs of the golfers in that event and the hole's par value and distance, EAGLE expected the group to average 0.145 strokes below par on number 14. But on the first day of that tournament, they actually combined for 0.435 strokes above par. That proved to be a reliable harbinger of its difficulty: the field did not come close to the model's rank-derived forecast on any of the three subsequent days. In the second round, they averaged 0.474 shots worse than par. And even after the weaker golfers were cut at the tournament's halfway mark, they could not do better than 0.401 above par in the third and fourth rounds.

Of course, Pebble Beach's number 14 is famous for a reason (and it is [now being](#)

modified to make it less brutal). But in one sense, it is highly typical: most holes that show such large discrepancies from EAGLE's base forecast early in a tournament tend to sustain them in later rounds. We decided we could take advantage of this trend by measuring it using a non-linear regression, to determine just how much of the effect we can expect to persist from day to day. After calculating this factor, we add it into EAGLE's forecast.

Next, EAGLE seeks to distinguish between golfers of equivalent OWGRs. Most importantly, not all world number ones are created equal. At the height of his powers, Tiger Woods won seven majors in three years. In contrast, Mr Day, today's top-ranked player, has just a single major title to his name. Moreover, the OWGRs suffer from other flaws. They ignore all data over two years old. And they are based exclusively on a player's finishes in tournaments rather than his individual scores. That means that a golfer's rank will get worse if he has the best round of his life on a day when a few competitors happen to do a little bit better. So our next step was to conduct another non-linear regression, to determine the optimal blend of a player's OWGR-based forecast alongside his own historical difficulty-adjusted scores in major tournaments.

Players outside the top 200 do not have a specific OWGR-based projection. As a result, we had to use the average performance of all unranked golfers in its stead. Unsurprisingly, this figure turned out to be quite a weak predictor. So unranked players, unless they have a large amount of recent major-tournament experience, tend to have very volatile projections: a few birdies in a row on the first day of a tournament are enough to increase EAGLE's estimate of their skill level significantly. In contrast, the OWGRs are impressively reliable. Even for the most frequently observed golfer in our dataset—Mr Scott, who shot almost 4,000 holes at major PGA tournaments between 2001 and 2015—we wound up weighting the OWGRs about twice as much as his personal record.

Max Planck's favourite sport

Thanks to these corrections, EAGLE can replace its simple projection based on OWGR, par and distance with a customised forecast that incorporates each hole's specific difficulty and each player's past performance. However, all of these calculations still rely on a third false assumption: that golf scores come in infinitely divisible pieces. The output of the equations outlined above is a projection of a player's expected score on a given hole—as good as 0.689 strokes below par, if you're Mr McIlroy on an impossibly short 284-yard par-four at the 2015 US Open (he eagled it), or as bad as 1.045 strokes above it if you had the misfortune to be a slumping Charles Challen on the eighth hole of the second round of the 2003 British Open (he bogeyed). In reality, no one gets a -0.689 or a +1.045, because there are no fractional strokes. There are only eagles, birdies, pars and the ever-worsening parade of bogey types—leaving aside the astonishingly rare three-under-par albatross, which occurred just eight times in our entire dataset.

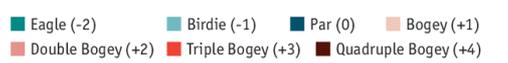
In order to translate these continuous projections into discrete scores, we used a cumulative ordered logit regression. This technique generates probabilities for seven different outcomes for each golfer on each hole, ranging from eagle to quadruple bogey. (Scores below -2 or above +4 were so rare that we could not model them reliably.) The resulting forecasts vary based on the hole's par value and distance, how the field performed on the hole previously and the golfer's skill. Chart 4 illustrates the range of potential outcomes. For the shortest and longest holes in our dataset of each par value, it lists the chances of each score for the best and worst golfers that played them.

Yes, 2 billion simulated holes

Finally, with these probabilities in hand, we deployed the Monte Carlo simulation method, originally developed by nuclear-weapons scientists in the 1940s. For each golfer on each

Twelve ways to slice the pie

Best and worst performers, selected holes, men's major golf tournaments 2001-16



Tournament	Golfer	Round	Hole number	Par	Distance Yards	Projected score to par	Share of probability, %	Actual Result
US Open 2013	BEST Tiger Woods	3	13	3	98	-0.18	PAR	Par
	WORST Michael Weaver	3	13	3	98	-0.06		Birdie
US Open 2016	BEST Jason Day	4	8	3	299	0.42		Par
	WORST Matt Marshall	4	8	3	299	0.60		Bogey
US Open 2013	BEST Tiger Woods	3	10	4	280	-0.25		Bogey
	WORST Matt Weibring	3	10	4	280	-0.03		Birdie
US Open 2015	BEST Rory McIlroy	2	13	4	551	0.18		Bogey
	WORST Alex Kim	2	13	4	551	0.70		Bogey
Masters 2001	BEST Tiger Woods	4	13	5	485	-0.48		Birdie
	WORST Billy Casper	2	13	5	485	0.07		Par
US Open 2016	BEST Jason Day	1	12	5	684	0.12		Par
	WORST Zach Edmondson	1	12	5	684	0.75		Bogey

Source: The Economist

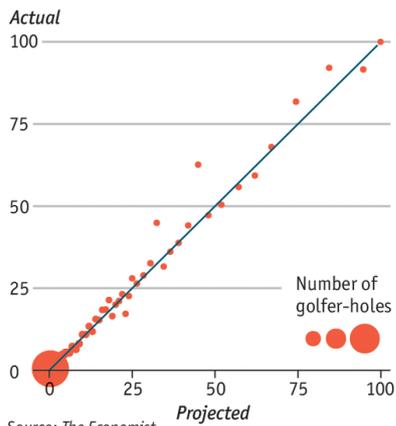
Economist.com

hole in each tournament, we pulled random numbers ranging from zero to one 10,000 times, using the previously-calculated probabilities to assign each number to an outcome like birdie or bogey. We then repeated the exercise as tournaments proceeded, simulating the remainder of the event after every golfer had completed the first hole, and again after they had completed the first and second, and so on until there was only one hole left to simulate. All told, we simulated over 2 billion golfer-holes. In each run, we allowed each player's projected hole outcomes to be influenced by earlier ones, so that a golfer who randomly happened to draw a few double bogeys to start a tournament would see his forecasts for the rest of that simulation deteriorate accordingly. Once the CPUs finished grinding, all we had to do was tally up the results and see how often each player won.

The resulting probabilities seem fairly well-calibrated (see chart 5). There were 351,000 golfer-holes where EAGLE pegged a player's likelihood of winning at less than 0.5%. It expected 0.1% of that group to defy the odds and win. Sure enough, precisely 0.1% did. (The most unlikely comeback was the Korean journeyman Y.E. Yang's astonishing rally in the 2009 PGA Championship after trailing Mr Woods by eight strokes, an event so unlikely it did not occur even once among our 10,000 simulations. The next-most-improbable was Keegan Bradley's victory at the 2011 PGA, when he fell seven strokes behind the leader in the first eight holes and wound up winning, a one-in-10,000 shot.) Similarly, players projected to win one time out of 20 did so 5% of the time, those forecast for one in ten prevailed 11% of the time and those predicted at one in five came through 20% of the time.

Right on target

EAGLE win probability vs actual, major men's golf tournaments 2001-16, %



Economist.com

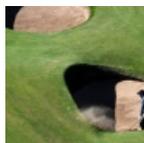
Because the vast majority of golfers go on to lose, there are far fewer cases when one has a strong chance of winning. As a result, the numbers at the high end are somewhat less reliable. For example, the 156 players estimated for a joint 45% probability actually won 63% of the time, and another 274 given a collective 79% forecast went on to win 86% of the time. But these are outliers: the predicted chances for most of the groups of

favourites were still spot-on. And there was no apparent pattern to the errors, suggesting they are likely to be mere consequences of smaller sample sizes rather than a systematic flaw.

I have to admit it's getting better

Moreover, EAGLE is a work in progress, and will improve with time. Above all, we still have not managed to rid the model of its unproven assumptions. Its Monte Carlo approach rests on the premise that each competitor's performance is unaffected by that of any other. In reality, golfers face strategic incentives to alter their play. One nursing a healthy lead should lay up and settle for par to avoid a Spieth-like collapse, whereas one trailing by a large margin would do well to take risks. Players may also be affected psychologically by their position in the standings. Similarly, EAGLE currently assumes that all golfers of equivalent skill will perform equally on holes of equivalent difficulty. This is almost surely false. There must be some degree of "fit", good or bad, between a golfer's individual strengths and weaknesses and the course he is playing. Future iterations of EAGLE will attempt to incorporate these dynamics.

In addition, EAGLE now weights a player's past performances based exclusively on how long ago they occurred. This could turn out to be an oversimplification. We'd like to explore the notion of "streakiness" or a "hot hand", and find out whether a player who happens to put together a few remarkable or terrible rounds in a row is likely to continue those trends in the short term. And finally, there are countless additional data sources that can enrich EAGLE. We plan to test the impact of tee times and weather, to separate out putting and driving performance, to use non-scorecard information like driving distances and accuracy and to expand our dataset to include non-major tournaments. Keep checking in at economist.com/eagle, both throughout the 2016 Open Championship and in the tournaments to come.

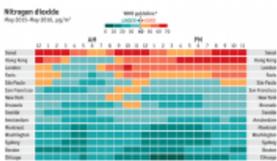
	<p>Previous</p> <p>Golf and politics: How Brexit might affect the Open Championship</p>	<p>Next</p> <p>Ranking performances in golf: Phil Mickelson is golf's best loser ever</p>	
--	---	---	---

Like 150 Tweet Share 53 G+ 4

[View all comments \(8\)](#)

[Add your comment](#)

More from the Economist



Daily chart: Comparing urban air pollution



Cinema absurdité: The surreal, dreamlike films of A24



The Economist explains: Why golf is back at the Olympics

- **Daily chart:** Rio's long decline
- **The Petrobras scandal:** A judge indicts Brazil's former president
- **The Economist explains:** How Clinton and Trump plan to boost wages
- **Following the flock:** Religious leaders grapple with doctor-assisted dying
- **Money talks:** Stressed out banks
- **Eastern European writing:** Svetlana Alexievich's monument to those left

behind

Classified ads



Contact us

Help

My account

Subscribe

Print edition

Digital editions

Events

Jobs.Economist.com

Timekeeper saved articles

Sections

- United States
- Britain
- Europe
- China
- Asia
- Americas
- Middle East & Africa
- International
- Business & finance
- Economics
- Markets & data
- Science & technology
- Special reports
- Culture
- Multimedia library

Debate and discussion

- The Economist debates
- Letters to the editor
- The Economist Quiz

Blogs

- Bagehot's notebook
- Buttonwood's notebook
- Democracy in America
- Erasmus
- Free exchange
- Game theory
- Graphic detail
- Gulliver
- Prospero
- The Economist explains

Research and insights

- Topics
- Economics A-Z
- Style guide
- The World in 2016
- Which MBA?
- MBA Services
- The Economist GMAT Tutor
- The Economist GRE Tutor
- Executive Education Navigator
- Reprints and permissions

The Economist Group »

- The Economist Intelligence Unit
- The Economist Intelligence Unit Store
- The Economist Corporate Network
- Ideas People Media
- 1843 Magazine
- Roll Call
- CQ
- EuroFinance
- The Economist Store
- Editorial Staff
- Modern Slavery Statement

[View complete site index »](#)

[Contact us](#) [Help](#) [About us](#) [Advertise with us](#) [Staff Books](#) [Careers](#) [Site index](#)

Copyright © The Economist Newspaper Limited 2016. All rights reserved. [Accessibility](#) [Privacy policy](#) [Cookies info](#) [Terms of use](#)