LETTER

An improved limit on the charge of antihydrogen from stochastic acceleration

M. Ahmadi[1], M. Baquero-Ruiz[2,3], W. Bertsche[4,5], E. Butler[6,7], A. Capra[8], C. Carruth[2], C. L. Cesar[9], M. Charlton[10], A. E. Charman[2], S. Eriksson[10], L. T. Evans[2], N. Evetts[11], J. Fajans[2], T. Friesen[12], M. C. Fujiwara[13], D. R. Gill[13], A. Gutierrez[11], J. S. Hangst[12], W. N. Hardy[11], M. E. Hayden[14], C. A. Isaac[10], A. Ishida[7], S. A. Jones[10], S. Jonsell[15], L. Kurchaninov[13], N. Madsen[10], D. Maxwell[10], J. T. K. McKenna[13], S. Menary[8], J. M. Michan[13], T. Momose[16], J. J. Munich[14], P. Nolan[1], K. Olchanski[13], A. Olin[13,17], A. Povilus[2], P. Pusa[1], C. Ø. Rasmussen[12], F. Robicheaux[18], R. L. Sacramento[9], M. Sameed[10], E. Sarid[19], D. M. Silveira[9], C. So[2], T. D. Tharp[12], R. I. Thompson[20], D. P. van der Werf[10], J. S. Wurtele[2,21] & A. I. Zhmoginov[2]

Antimatter continues to intrigue physicists because of its apparent absence in the observable Universe. Current theory requires that matter and antimatter appeared in equal quantities after the Big Bang, but the Standard Model of particle physics offers no quantitative explanation for the apparent disappearance of half the Universe. It has recently become possible to study trapped atoms[1–4] of antihydrogen to search for possible, as yet unobserved, differences in the physical behaviour of matter and antimatter. Here we consider the charge neutrality of the antihydrogen atom. By applying stochastic acceleration to trapped antihydrogen atoms, we determine an experimental bound on the antihydrogen charge, Qe, of $|Q| < 0.71$ parts per billion (one standard deviation), in which e is the elementary charge. This bound is a factor of 20 less than that determined from the best previous measurement[5] of the antihydrogen charge. The electrical charge of atoms and molecules of normal matter is known[6] to be no greater than about $10^{-21}e$ for a diverse range of species including H_2, He and SF_6. Charge–parity–time symmetry and quantum anomaly cancellation[7] demand that the charge of antihydrogen be similarly small. Thus, our measurement constitutes an improved limit and a test of fundamental aspects of the Standard Model. If we assume charge superposition and use the best measured value of the antiproton charge[8], then we can place a new limit on the positron charge anomaly (the relative difference between the positron and elementary charge) of about one part per billion (one standard deviation), a 25-fold reduction compared to the current best measurement[8,9].

The charge of the antihydrogen atom can be inferred from the charge anomalies of the antiproton and positron. The charge anomaly of the antiproton is known from spectroscopic measurements[8,10] on antiprotonic helium, $\overline{p}He^+$, and is $||q_{\overline{p}}| - e|/e < 0.7$ parts per billion (p.p.b.) with a confidence level of 90%, where $q_{\overline{p}}$ is the antiproton charge. The charge anomaly of the positron[8,9,11] is less well known: $|q_{e^+} - e|/e < 25\,\text{p.p.b.}$, where q_{e^+} is the positron charge (no confidence level given, but assumed here to be 1σ), determined by measurements of the positron cyclotron frequency and the positronium Rydberg constant. Assuming charge superposition, we can therefore combine these two anomalies to infer an indirect experimental bound of $|Q| \lesssim 25\,\text{p.p.b.}$ on the charge Qe of antihydrogen. This bound is much looser than the

direct experimental bounds for normal-matter atoms; however, the methods used to set these normal-matter bounds generally use macroscopic quantities of atoms and so are unsuitable for studies of antihydrogen, of which only approximately 1,000 antiatoms have been trapped and detected[1–4].

The earliest direct experimental limit on the charge of antihydrogen, $|Q| < 0.02$, came from the observation[12] that Lorentz forces were insufficient to deflect an energetic (momentum of $1.94\,\text{GeV}\,c^{-1}$) beam[13] of antihydrogen atoms away from a detector. A much more precise bound, $Q = (-13 \pm 11 \pm 4)\,\text{p.p.b.}$ (1σ confidence level), was recently determined[5] by searching for a static-electric-field-induced shift in the average position of antiatoms confined in the ALPHA-I antihydrogen trap at CERN. This shift was measured by recording the position along the electric-field axis as the antiatoms were released from the trap by turning off the trapping magnetic fields.

The precision of this static shift measurement was limited by the magnitude of the applied electric fields, and by statistical and systematic issues related to accurately measuring small deflections. The sensitivity of the charge measurement can be enhanced by using stochastic (that is, randomly time-varying) electric fields[5,14,15]. Under the influence of such fields, a putatively charged, trapped antihydrogen atom would experience small, stochastic energy kicks. On average, the antiatom would gradually gain enough energy from these kicks to escape the shallow potential well in which it was trapped. This process can be modelled as a random walk similar to Fermi acceleration[16]. By searching for the loss of antiatoms caused by these stochastic electric fields, we set a new limit on $|Q|$.

Our measurement was performed on antihydrogen atoms confined in the ALPHA-2 trap at CERN. This trap uses a minimum-magnetic-field (minimum-B) configuration to trap the antiatoms. Schematics of the trap and fields are shown in Fig. 1a and b, respectively. The magnetic minimum is formed by mirror and octupole magnetic coils. Antihydrogen atoms whose magnetic moments are oriented antiparallel to the trap magnetic fields are attracted to the magnetic minimum at the trap centre; for our fields, which are of the order of 1 T, the depth of the antihydrogen potential well is $E_{well} = 0.54\,\text{K}$, where E_{well} is measured in temperature units.

The general methods by which we trap antihydrogen are described elsewhere[1–3,17], but briefly, antihydrogen atoms are created by first trapping separate cold plasmas containing several million positrons
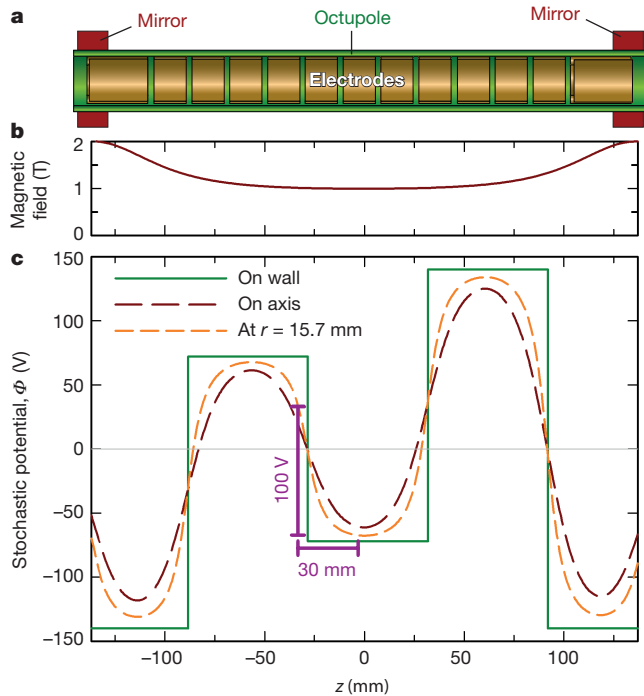
Figure 1 | Experiment schematic. a, Electrode structure of the ALPHA-2 trap, showing the approximate positions of the mirror and octupole coils. The particle detector (not shown) surrounds the octupole. (The vertical scale has been expanded for clarity.) **b**, Magnetic field on the axis (z) of the atom trap showing the axial magnetic potential well induced by the mirror fields. **c**, A snapshot of the applied stochastic electric-field potential Φ at several radii r as a function of z. 'On wall' indicates the wall of the Penning trap; 'on axis' indicates the axis of the trap ($r = 0$). The purple brackets indicate a typical travel distance and potential change experienced by an antiatom with 0.1 K of axial energy in 1 ms.

and several tens of thousands of antiprotons. The positrons come from a radioactive sodium source[18], and the antiprotons are generated by CERN's Antiproton Decelerator (AD). Antiatoms are formed by three-body recombination when these two plasmas are mixed using autoresonance[19,20] in the trap interior. Approximately 0.01% of the roughly ten thousand atoms created in each 1-s-long mixing interval are sufficiently cold to be trapped. The antiatoms are detected by releasing them from the trap by turning off the trap magnetic fields. A three-layer silicon vertex detector (SVD) surrounding the trap is used to detect, locate and time the resulting antiatom annihilations[21]. The detector and its associated multivariate analysis[22–24] (MVA) software can detect approximately 82% of the antiatoms that annihilate in the trap.

To generate the stochastic electric field necessary for our measurement, we impose the electric potential shown in Fig. 1c on the trap by biasing the trap electrodes. We make this potential, and the associated electric field, vary stochastically by repeatedly changing the bias on the electrodes from that shown in Fig. 1c to its negation. We use $N = 84,900$ transitions, which, altogether, last 114.9 s in each experimental cycle (see Fig. 2). The interval between transitions averages 1 ms, but is randomized for reasons that are discussed later.

Each of the potential transitions would non-adiabatically perturb the kinetic energy of a charged antiatom; as noted above, these kicks correspond, approximately, to a random walk in energy. To model this process, we need to estimate the average potential change subsequent to each of the transitions. Measurements using our first trap, ALPHA-I[3], showed that the energy of the trapped antiatoms is roughly distributed[2,5,25,26] as a high-temperature, three-dimensional (3D) Maxwellian distribution truncated near E_{well}. The typical energy of the component of motion along the trap axis is therefore about 0.1 K, which corresponds to a travel distance of about 30 mm in 1 ms. From the purple
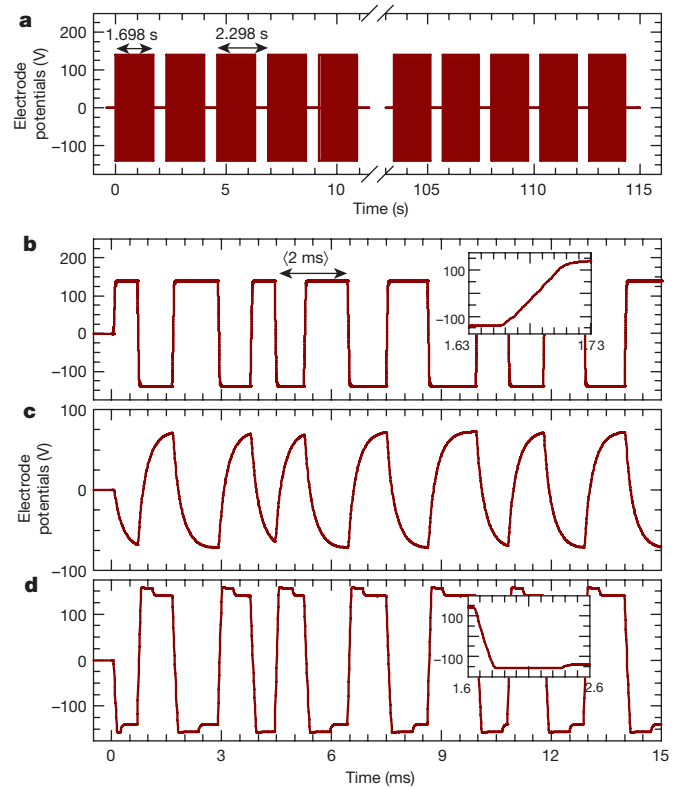

Figure 2 | Stochastic drive potentials. a, Time structure of the stochastic drive potential repetitions. The potential is non-zero within the first 1.698 s of each repetition. The fifty repetitions are separated by 2.298 s. **b–d**, Snapshots of the typical time histories of the measured potentials applied to different electrodes; insets illustrate details of the transitions. Note that the time intervals between transitions vary stochastically.

brackets shown in Fig. 1c, we see that an antiatom possessing this axial kinetic energy would experience a potential change of $\Delta\Phi \approx 100$ V in 1 ms, although there are substantial uncertainties in this estimate due to variations in the potential over the volume of the trap, the distribution of kinetic energies and the potential changes arising from radial motion. A simple energy-diffusion model predicts that an antiatom with a putative charge Qe would gain energy of about $|Q|e\Delta\Phi\sqrt{N}$. Consequently, the antiatom would gain roughly enough energy to escape from a trap of depth E_{well} if

$$|Q| \gtrsim \frac{E_{well}}{e\Delta\Phi\sqrt{N}} \qquad (1)$$

Evaluating this equation using our experimental parameters yields $|Q| \gtrsim 1.6$ p.p.b. However, equation (1) assumes that the antiatom starts with zero initial kinetic energy. Prior measurements indicate that the average antiatom starts with more than half the energy it needs to escape; thus, the bound on $|Q|$ is probably lower.

Stochastic acceleration requires an element of randomness in the relation between the frequency of the driving-potential inversions and the frequencies of the antiatom orbital oscillations. Trajectory simulations (see, for example, figure 6 in ref. 15) demonstrate that the orbits in the unforced system are generally not regular and do not have a fixed frequency. In many cases, these variations in orbital frequency would make even a strictly periodic driving potential act stochastically. However, some orbits have long-lasting periods of regular motion, and antiatoms in these orbits might not be accelerated by the driving potential. As proposed in ref. 15, we introduce sufficient stochasticity into the system by a random modulation, with a uniform distribution and a standard deviation of 0.2 ms, of the time interval between potential transitions. (This variability is visible in Fig. 2b–d.) Simulations[14,15] show that it is not necessary

Table 1 | Results of the stochastic and null trials

	Number of trials	Observed antiatoms
Stochastic trials	10	12
Null trials	10	12

or beneficial to randomize the voltage levels between which the electrode potentials switch.

To bound $|Q|$, we alternated between two experimental protocols, each with four phases. The first protocol involved stochastic trials, in which we (1) trapped antihydrogen, (2) subjected any trapped antihydrogen to the stochastic potential variations described above, (3) allowed any antiatoms remaining after the stochastic phase to escape by turning the confining magnetic fields off, and then (4) counted the number of antiatoms that thus escaped using the silicon vertex detector. The second protocol involved null trials, which were performed in strict alternation with the stochastic trials, and were identical to them in all respects, except that the potentials were set to zero during phase (2). Thus, in both types of trials we attempted to hold the antiatoms for the same time (114.9 s) after phase (1). Assuming that the underlying initial antihydrogen trapping rate and all other systematic effects are the same in the stochastic and null trials, a statistically significant difference between the number of antiatoms detected after these trials would imply the existence of a non-zero Q being acted upon by the oscillating electric fields of the stochastic trials. (As discussed in Methods, known polarization and polarizability effects are negligible in the current experiment.)

The results of executing both protocols ten times are presented in Table 1. We observed no differences in the number of antiatoms released in phase (3) between the two protocols, which suggests that the charge—if any—of antihydrogen is below that required to escape from the trap, $|Q| \lesssim 1.6$ p.p.b. (based on the analytic estimate given above). This rough estimate can be made more robust and precise by simulating the dynamics of the antiatoms and using a proper statistical analysis. (Events detected during phase (2) are compiled in Extended Data Table 1 and discussed in Methods, and are compatible with cosmic rays.)

We use simulations[1,2,5,15,25,26] that model the experiment closely, including detailed models of the trap fields and stochastic timing, and realistic initial conditions validated by previous experiments. The simulations begin by mimicking the initial mixing of the two plasmas and subsequent relaxation periods. We model the null trials by propagating the antiatoms in the electric-field-free trap for a further 114.9 s. We model the stochastic trials by assigning the antiatoms a putative Q, and propagating them in the presence of the stochastic electric fields for this same 114.9 s. For each value of Q studied, we launch 1,000 antiatoms.

For each simulated Q, we determine the survival probability s by calculating the ratio of the number of antiatoms that survive the stochastic simulations to the number that survive the null simulations; the results are shown in Fig. 3. Because the stochastic drive reverses direction every half-cycle, it is unbiased on average; consequently, s depends on the magnitude of the charge $|Q|$, but not on its sign.

Although it is clear from Fig. 3 that the bound on $|Q|$ set by our experiment will be in the neighbourhood of 1 p.p.b., it is not immediately clear exactly what s and, hence, what $|Q|$ to assign from this figure and the data in Table 1. We might begin by setting a threshold of $s = 0.5$; however with a per-trial antiatom observation rate of $r \approx 24/20 = 1.2$, the probability that 12 (what we observed) or more antiatoms would have survived in 10 trials with $s = 0.5$ is only 0.02 assuming Poissonian statistics. Clearly, our data will support setting the threshold at some larger value of s. Furthermore, our limited data set only provides an estimate for r. This suggests that a Bayesian methodology, in which r is treated as a nuisance parameter, would be useful for finding a more precise limit on $|Q|$.

In Methods we present a simple Bayesian analysis that suggests that $s > 0.79$ and that the corresponding bound on $|Q|$, found by inverting the curve in Fig. 3, is $|Q| < 0.59$ p.p.b. However, this analysis ignores
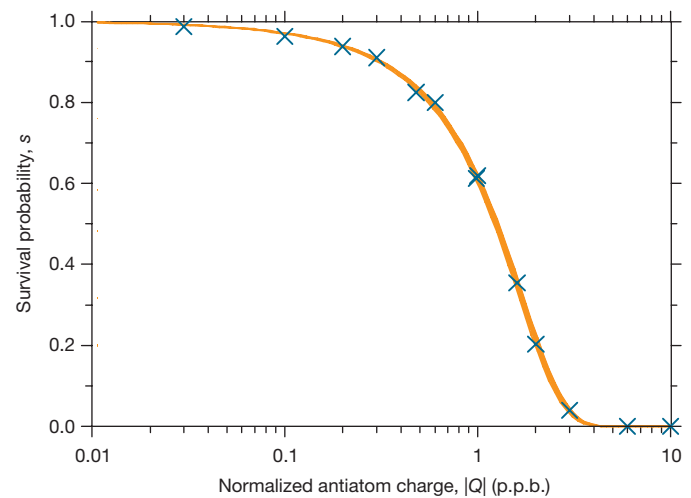

Figure 3 | Survival probability. Simulated survival probability s as a function of $|Q|$ for the stochastic trials. The blue crosses are the number of antiatoms surviving at the given $|Q|$ value divided by the number of antiatoms surviving at $Q = 0$ (the null simulation). The orange band of varying thickness is the 1σ confidence region from a Bayesian fit to this simulation data; see the Supplementary Information for further explanation.

systematic uncertainties; fortunately, our experiment is relatively insensitive to these uncertainties. For example, the analytic expression given in equation (1) shows that uncertainties in Φ (or in the magnitude of the electric field $|E|$) affect our bound on $|Q|$ linearly. Because these fields are known to within approximately 1%, uncertainties in Φ have a negligible effect compared to that of our statistical uncertainties. Similarly, uncertainties in the magnitude of the magnetic field $|B|$ contribute to uncertainties in $|Q|$ through its linear dependency on E_{well}. Because E_{well} depends on the mirror and octupole fields approximately linearly, and because both are known to within approximately 1%, these effects are unimportant. Details of the magnetic field structure beyond the maximum and minimum are similarly unimportant because they do not affect the trap depth. Simulations show that s is not greatly influenced by changes in the volume from which the antiatoms are initially launched.

All of the data were taken over adjacent half-shifts, separated by 16 h, at CERN's AD facility. We are not aware of any mechanism that could have substantially altered the observation rate r in synchronization with the strict alternation of the stochastic and null trials. Likewise, this alternation makes the measurement robust against failures of the detector system. Nonetheless, we performed extensive tests[27] of the detector performance to confirm that there were no substantial failures or performance degradations (see Extended Data Fig. 1). The expected number of cosmic rays incorrectly identified as antiatoms is 0.07 for each of the two sets of ten trials; this has a negligible effect on our bound on $|Q|$ (see Supplementary Information).

Simulations do reveal a small dependence of s on the initial energy distribution of the antiatoms. Two distributions that bound our preferred, truncated Maxwellian distribution were studied for our previous apparatus ALPHA-I. Although we have good evidence to rule out these bounding distributions[5,25,26], their impact on the measurement reported here is discussed in Supplementary Information. With an additional, even smaller correction arising from the statistical error in the Monte Carlo simulation, our bound degrades to $|Q| < 0.71$ p.p.b. at a 68.3% (1σ) confidence level. This bound represents a 20-fold improvement on the best previous bound[5].

Given that the charge anomaly of the antiproton[8,10] is bound at below 0.7 p.p.b., our result, combined with the assumption of charge superposition, limits the charge anomaly of the positron to 1 p.p.b. (1σ), a 25-fold improvement on the best previous bound[8,9]. By modestly increasing the potentials, holding the antiatoms for longer times and cooling the antiatoms[28], we expect (see the scalings in equation (1)) to

be able to improve the antihydrogen charge bound by several orders of magnitude[15]; these realizable improvements exceed those that could be obtained by the previously used static-electric-field shift methodology[5].

1. Andresen, G. B. et al. Trapped antihydrogen. Nature **468**, 673–676 (2010).
2. Andresen, G. B. et al. Confinement of antihydrogen for 1,000 seconds. Nature Phys. **7**, 558–564 (2011).
3. Amole, C. et al. The ALPHA antihydrogen trapping apparatus. Nucl. Instrum. Methods Phys. Res. A **735**, 319–340 (2014).
4. Gabrielse, G. et al. Trapped antihydrogen in its ground state. Phys. Rev. Lett. **108**, 113002 (2012).
5. Amole, C. et al. Experimental limit on the charge of antihydrogen. Nature Commun. **5**, 3955 (2014).
6. Bressi, G. et al. Testing the neutrality of matter by acoustic means in a spherical resonator. Phys. Rev. A **83**, 052101 (2011).
7. Quigg, C. Gauge Theories of the Strong, Weak, and Electromagnetic Interactions (Westview Press, 1997).
8. Olive, K. A. et al. Review of particle physics. Chinese Phys. C **38**, 090001 (2014).
9. Fee, M. S. et al. Measurement of the positronium $1\ ^3S_1$–$2\ ^3S_1$ interval by continuous-wave two-photon excitation. Phys. Rev. A **48**, 192–219 (1993).
10. Hori, M. et al. Two-photon laser spectroscopy of antiprotonic helium and the antiproton-to-electron mass ratio. Nature **475**, 484–488 (2011).
11. Hughes, R. J. & Deutch, B. I. Electric charges of positrons and antiprotons. Phys. Rev. Lett. **69**, 578–581 (1992).
12. Greenland, P. T. Antimatter. Contemp. Phys. **38**, 181–203 (1997).
13. Baur, G. et al. Production of antihydrogen. Phys. Lett. B **368**, 251–258 (1996).
14. Baquero-Ruiz, M. Studies on the Neutrality of Antihydrogen. PhD thesis, University of California, Berkeley (2013).
15. Baquero-Ruiz, M. et al. Using stochastic acceleration to place experimental limits on the charge of antihydrogen. New J. Phys. **16**, 083013 (2014).
16. Fermi, E. On the origin of the cosmic radiation. Phys. Rev. **75**, 1169–1174 (1949).
17. Andresen, G. B. et al. Search for trapped antihydrogen. Phys. Lett. B **695**, 95–104 (2011).
18. Murphy, T. & Surko, C. Positron trapping in an electrostatic well by inelastic collisions with nitrogen molecules. Phys. Rev. A **46**, 5696–5705 (1992).
19. Fajans, J. & Frièdland, L. Autoresonant (nonstationary) excitation of a pendulum, Plutinos, plasmas and other nonlinear oscillators. Am. J. Phys. **69**, 1096–1102 (2001).
20. Andresen, G. B. et al. Autoresonant excitation of antiproton plasmas. Phys. Rev. Lett. **106**, 025002 (2011).
21. Amole, C. et al. Silicon vertex detector upgrade in the ALPHA experiment. Nucl. Instrum. Methods Phys. Res. A **732**, 134–136 (2013).
22. Narsky, I. StatPatternRecognition: a C++ package for statistical analysis of high energy physics data. Preprint at http://arxiv.org/abs/physics/0507143 (2005).
23. Narsky, I. Optimization of signal significance by bagging decision trees. Preprint at http://arxiv.org/abs/physics/0507157 (2005).
24. Amole, C. et al. Resonant quantum transitions in trapped antihydrogen atoms. Nature **483**, 439–443 (2012).
25. Amole, C. et al. Discriminating between antihydrogen and mirror-trapped antiprotons in a minimum-B trap. New J. Phys. **14**, 015010 (2012).
26. Amole, C. et al. Description and first application of a new technique to measure the gravitational mass of antihydrogen. Nature Commun. **4**, 1785 (2013).
27. Capra, A. Testing CPT and antigravity with trapped antihydrogen at ALPHA. PhD thesis, York University (2015).
28. Donnan, P. H., Fujiwara, M. C. & Robicheaux, F. A proposal for laser cooling antihydrogen atoms. J. Phys. B **46**, 025302 (2013).

METHODS

Stochastic potentials. The stochastic potentials are applied by biasing the electrically isolated trap electrodes shown in Fig. 1a. Within the approximately 280-mm-long trapping region, we impose roughly 2.25 spatial periods of an electric potential that varies between approximately ±100 V (see Fig. 1c). These potentials come in fifty time-identical repetitions (see Fig. 2a). Individual repetitions last 1.698 s and are spaced at intervals of 2.298 s. Each repetition begins and ends with the potentials zeroed. Within each repetition, the potential cycles between that shown in Fig. 1c and its negation; the average time interval between transitions is 1 ms. Thus, there are 1,698 transitions in each repetition and $N = 84{,}900$ transitions in total. Typical time histories of the potentials are shown in Fig. 2b–d. The transitions begin in synchronization for all electrodes; the differences between the attacks (the rise times of each pulse) and magnitudes of the three potentials shown in Fig. 2b–d are due to deliberate differences in the amplifiers and filters driving the electrodes. These differences are precisely modelled in our simulations, and are not otherwise germane to the results presented here.

Data collection. All data were collected in two adjacent half-shifts, separated by an off-shift period, at CERN's AD. The data collection followed the protocol outlined above of strict alternation of stochastic and null trials. Data collection ended when we reached the predetermined number (ten) of trials of each type. Trials that had an immediately obvious failure (no antiprotons delivered from the AD, premature failure of our magnet system and so on) were aborted and repeated; otherwise, all the data that were collected are presented here.

Equation of motion. The antiatom equation of motion in the trap is

$$M\ddot{\boldsymbol{x}} = \boldsymbol{\nabla}(\boldsymbol{\mu}\cdot\boldsymbol{B}) + Qe(\boldsymbol{E}+\dot{\boldsymbol{x}}\times\boldsymbol{B}) \qquad (2)$$

Here, $\boldsymbol{x} = \boldsymbol{x}(t)$ is the centre-of-mass position of an antiatom at time t, $\boldsymbol{E} = \boldsymbol{E}(\boldsymbol{x},t)$ and $\boldsymbol{B} = \boldsymbol{B}(\boldsymbol{x})$ are the applied electric and magnetic fields, M is the total mass and $\boldsymbol{\mu}$ is the magnetic moment of antihydrogen. A trapped antiatom must be low-B-field-seeking, so the direction of $\boldsymbol{\mu}$ is assumed to adiabatically track the direction antiparallel to $\boldsymbol{B}(\boldsymbol{x}(t))$.

Simulations. The simulations mimic the experimental procedures closely, and are similar to those in refs 1, 2, 5, 14, 15, 25 and 26, in which numerous tests of robustness of the dynamical model and numerical approximations have been performed. Simulated antiatoms are launched from a volume corresponding to the precursor positrons. The antiatoms are drawn from an energy distribution obeying a truncated 3D Maxwellian, which scales as $E_0^{1/2}$ at low energies, where E_0 is the initial energy of the antiatoms. The truncation energy is set to 0.75 K, well above the trap depth of $E_{\text{well}} = 0.54$ K, but well below the positron temperature of tens of kelvin that sets the overall Maxwellian temperature. As discussed below, the antiatoms are then held for 1.2 s. Most of the antiatoms with energies above E_{well} will be lost during this relaxation/randomization time; however, a few, on quasi-bound orbits, will remain[2,29].

In the experiment, we mix positrons and antiprotons for 1 s. Antihydrogen production peaks at the beginning of the mixing phase, but we do not know when within the mixing phase the antihydrogen that is trapped is created. After mixing, the antiatoms are allowed to randomize for 0.4 s. During this relaxation/randomization phase, electric fields are applied that clear the positrons and antiprotons remaining in the trap. We do not model these fields because they are generally weak compared to the stochastic potentials, there are few transitions and, experimentally, both the stochastic and null trials experience these fields. Altogether, these phases take a total of 1.4 s; in the simulations, to account for the uncertainty in the synthesis time, we allot 1.2 s for this composite relaxation/randomization time. We note that 1.2 s is short compared to the 114.9-s stochastic/hold phase, and that relaxation continues during this phase. Thus, the 0.2-s uncertainty has a negligible effect on our results.

We launch 1,000 antiatoms for each case we study. About 67% of these antiatoms survive the initial relaxation phase. Because 61% have initial energy below E_{well}, the remaining 6% are on quasi-bound orbits.

As discussed in the main text, we then propagate the antiatoms remaining after the relaxation phase for 114.9 s without the stochastic electric fields (null trials), or with the stochastic fields and with an assumed Q (stochastic trials). At the end of the 114.9 s, we count the number of antiatoms that have survived the simulation to determine the survival probability; there is no need to model the magnetic field shut-off used to count the antiatoms in the experiment.

We use a symplectic leap-frog propagator to numerically integrate equation (2). The electric potentials in the simulation are determined from a precise model of the trap geometry, using the COMSOL finite element code[30] to solve the Laplace equation, and the measured time-history of the potentials on each electrode (see Fig. 2). The magnetic field is calculated from high-accuracy analytic expansions derived from numerical Biot–Savart modelling[25]. We do not see any interesting effects from coarsening the field solutions (see also ref. 5).

Simulation convergence studies show that the step size need not be shorter than 4 μs; we used 0.8 μs for the simulations reported here. They also indicate that the electrode drive amplifier response time need only be known to the 4-μs level, which is well within its measured accuracy.

Bayesian analysis. In the Supplementary Information, we present a complete Bayesian analysis that explicitly accounts for uncertainty in the relationship between s and $|Q|$ due to the finite number of samples in our Monte Carlo simulations, and for various systematic effects. Here, we give an abbreviated analysis. We search for the value of the survival probability, s^*, that yields a $(1-\alpha)$ posterior credible interval given our observations

$$1-\alpha = \int_{s^*}^{1}\mathrm{d}s'p(s'\,|\,N_{\text{s}}, N_{\text{n}}) \qquad (3)$$

Here, $p(s\,|\,N_{\text{s}}, N_{\text{n}})$ is the probability density for the survival probability s, given a total of N_{s} observed atoms in the stochastic trials and N_{n} observed atoms in the null trials, and $1-\alpha = 0.683$ for the equivalent of a 1σ credible interval. Bayesian formalism gives

$$p(s|N_{\text{s}}, N_{\text{n}}) = \frac{\int_{0}^{\infty}\mathrm{d}r'p(N_{\text{s}}, N_{\text{n}}\,|\,s, r')\pi(s)\pi(r')}{\int_{0}^{1}\mathrm{d}s'\int_{0}^{\infty}\mathrm{d}r'p(N_{\text{s}}, N_{\text{n}}\,|\,s', r')\pi(s')\pi(r')} \qquad (4)$$

where $\pi(s)$ and $\pi(r)$ are the priors on s and r, respectively. The likelihood $p(N_{\text{s}}, N_{\text{n}}\,|\,s, r)$ is given by

$$p(N_{\text{s}}, N_{\text{n}}\,|\,s, r) = \sum_{n\geq N_{\text{s}}}^{\infty}\left[\binom{n}{N_{\text{s}}}s^{N_{\text{s}}}(1-s)^{n-N_{\text{s}}}\right]\left[\frac{(Kr)^n}{n!}\mathrm{e}^{-Kr}\right]\left[\frac{(Kr)^{N_{\text{n}}}}{N_{\text{n}}!}\mathrm{e}^{-Kr}\right] \qquad (5)$$

$$= \frac{(Krs)^{N_{\text{s}}}}{N_{\text{s}}!}\mathrm{e}^{-Krs}\frac{(Kr)^{N_{\text{n}}}}{N_{\text{n}}!}\mathrm{e}^{-Kr} \qquad (6)$$

where $K = 10$ is the number of stochastic and null trials. The first factor in the summation in equation (5) is a binomial subsampling of the Poisson-distributed number (given by the second factor) of antiatoms n present in the stochastic trials. (The Poisson assumption is further discussed in the Supplementary Information.) This subsampling models loss, with probability $1-s$, due to the stochastic fields. The third factor is the Poisson-distributed number of antiatoms present in the null trials. The first two terms can be combined into a single Poisson process, yielding equation (6).

The bound on $|Q|$ compatible with the current experiment is substantially lower than the previously reported bound[5], and we do not possess accurate prior information about the actual observation rate r in the particular trials used here. Hence, we choose to use uninformative priors. Various notions of uninformativeness have been suggested[31]. For both binomial and Poisson sampling distributions, the widely used Bernardo–Berger reference priors[32], Jeffreys invariant priors[33] and Rodriguez entropic priors[34] all coincide, and are given by

$$\pi(s) \propto \frac{1}{\sqrt{s(1-s)}}\,;\quad \pi(r) \propto \frac{1}{\sqrt{r}} \qquad (7)$$

When evaluated, equations (3)–(7) suggest that $s > 0.79$ and that the corresponding bound on $|Q|$, found by inverting the curve in Fig. 3, is $|Q| < 0.59$ p.p.b. without considering systematic uncertainties.

Bounding the electric dipole moment. If we take $Q = 0$, but change the equation of motion, equation (2), to

$$M\ddot{\boldsymbol{x}} = \boldsymbol{\nabla}(\boldsymbol{\mu}\cdot\boldsymbol{B}) + \alpha\boldsymbol{\nabla}|\boldsymbol{E}|^2 \qquad (8)$$

to include the effects of electric polarizability, then we can reinterpret our experiment as setting a bound on the polarizability α of the antihydrogen atoms (in the absence of appreciable permanent monopole or dipole moments). This sets a limit of $\alpha < 3\times10^6\alpha_0$, at 68.3% confidence, not including systematic effects, where $\alpha_0 = 4\pi\varepsilon_0 \times 0.667\times10^{-24}$ cm^3 is the known polarizability of hydrogen[8]. Note that effects due to polarizability will ultimately limit the precision of the stochastic methodology.

If instead we take $Q = 0$ and $\alpha = 0$, and change the equation of motion to

$$M\ddot{\boldsymbol{x}} = \boldsymbol{\nabla}(\boldsymbol{\mu}\cdot\boldsymbol{B}) + \boldsymbol{\nabla}(\boldsymbol{p}\cdot\boldsymbol{E}) \qquad (9)$$

then we can interpret our experiment as setting a bound on a permanent electric dipole moment $\boldsymbol{p}$ of antihydrogen. This gives a bound of $10^{-9}e$ cm at 68.3% confidence not including systematic effects. If this moment were to come from, for instance, a positron dipole moment, then quantum corrections[35] would loosen

the bound to about $10^{-5}e$ cm. This bound is many orders of magnitude greater than the current bound on the electron dipole moment of $9.7 \times 10^{-29}e$ cm (ref. 36).

Identification of antihydrogen. Antihydrogen detection, and cosmic ray rejection, is based on previous work within the ALPHA experiment (see ref. 37). By taking the energy deposited in each strip of the SVD, clusters of activated strips can be summed into 'raw hits'; raw hits from either side of a silicon module can be paired to make 'hits'; hits can be compared in 3D space to create 'tracks'; and tracks can be combined to point to an annihilation vertex.

The detector only detects the pions that are produced by the antiproton annihilations when antihydrogen atoms annihilate on the trap wall or on the very-low-density background gas; the detector is insensitive to the γ-rays produced by the simultaneous positron annihilations. We have developed many techniques for first eliminating antiprotons from our trap before the detection phase, and then discriminating between antihydrogen annihilations and the unlikely antiproton annihilations that might occur if some unbound antiprotons were still present[1,17,25]. We are confident that the events reported here are due to antihydrogen annihilations only.

To determine the best significance of antihydrogen detection against background, we fed multiple selection variables that measure the event shape and orientation into a MVA[22–24]. Complete reconstruction of the annihilation vertex is only part of the input to this MVA, because this method combines low- and high-level reconstruction data in a machine-learning algorithm. We used a large annihilation data sample and a large, independent, background data sample to train the machine-learning algorithm. The annihilation samples have fewer than 1% cosmic events and the background sample consists entirely of background events because the data are collected in the absence of antiprotons in the experiment. We then tested the performance of the MVA technique on a dedicated, separate test sample of data. Once trained, we used the MVA to analyse events that were detected during the experimental trials, and classified each event as either an antihydrogen annihilation event or as a background cosmic event. Note that we can tune the MVA differently for significance optimization during the stochastic/hold phase (2) and the release phase (3); for instance, we can decrease the false-negative rate where possible (increase the detection efficiency), at the cost of increasing the false-positive rate.

MVA of the release phase (3). With the MVA optimized for the 30-ms observation window of the release phase (3), $(82.3 \pm 0.2)\%$ of the annihilation events are accepted with a background rate of $242\,\text{mHz} \pm 2\%$, which gives $(0.073 \pm 0.002) \ll 1$ expected cosmic rays in the observation window of both the null and stochastic trials. The results of these trials are reported in Table 1, and are the primary data used here.

MVA of the stochastic/hold phase (2). For the 10 null and 10 stochastic trials, events detected during the stochastic/hold phase (2) (before release) were fed into a differently trained MVA classifier that has a detector acceptance of $(46.1 \pm 0.3)\%$ and a cosmic background rate of 6 mHz with 11% uncertainty (see Extended Data Table 1). Stochastic/hold phase events might have arisen from the escape of quasi-bound antiatoms, from annihilations on residual gas, or from ejection of the antiatoms by the stochastic acceleration. We performed a detailed statistical analysis to determine whether anything other than background needs to be introduced to account for the observed counts in Extended Data Table 1. That analysis (see Supplementary Information) shows that there is no reason to reject the assumption that these phase (2) counts are due to background, and thus do not provide reliable additional information regarding Q.

Detector performance. Because consistency between the detector performance during the null and stochastic trials is critical for this analysis, we rigorously scrutinized and cross-checked the performance of the SVD with scintillation detectors placed on the outside of the apparatus. We examined several quality-of-data (QOD) criteria using both cosmic data from dedicated cosmic observation windows, and annihilation data from the mixing phase of experimental trials. We found no anomalous differences between the stochastic and null trials. The results of these tests[27] are summarized in Extended Data Fig. 1.

29. Bowman, J. D. & Penttila, S. I. On the measurement of the neutron lifetime using ultracold neutrons in a vacuum quadrupole trap. *J. Res. Natl Inst. Stand. Technol.* **110,** 361–366 (2005).
30. COMSOL, http://www.comsol.com/ (COMSOL, Inc., 2015).
31. Kass, R. E. & Wasserman, L. The selection of prior distributions by formal rules. *J. Am. Stat. Assoc.* **91,** 1343–1370 (1996).
32. Berger, J. O. *Statistical Decision Theory and Bayesian Analysis* 2nd edn (Springer, 1985).
33. Jeffreys, H. *Theory of Probability* 3rd edn (Clarendon Press, 1961).
34. Rodriguez, C. C. in *Maximum Entropy and Bayesian Methods* (ed. Fougère, P. F.) Vol. 39 of *Fundamental Theories of Physics* 31–39 (Springer, 1990).
35. Sandars, P. G. H. The electric-dipole moments of an atom II. The contribution from an electric-dipole moment on the electron with particular reference to the hydrogen atom. *J. Phys. B* **1,** 511–520 (1968).
36. The ACME Collaboration. Order of magnitude smaller limit on the electric dipole moment of the electron. *Science* **343,** 269–272 (2014).
37. Andresen, G. B. *et al.* Antihydrogen annihilation reconstruction with the ALPHA silicon detector. *Nucl. Instrum. Methods Phys. Res. A* **684,** 73–81 (2012).
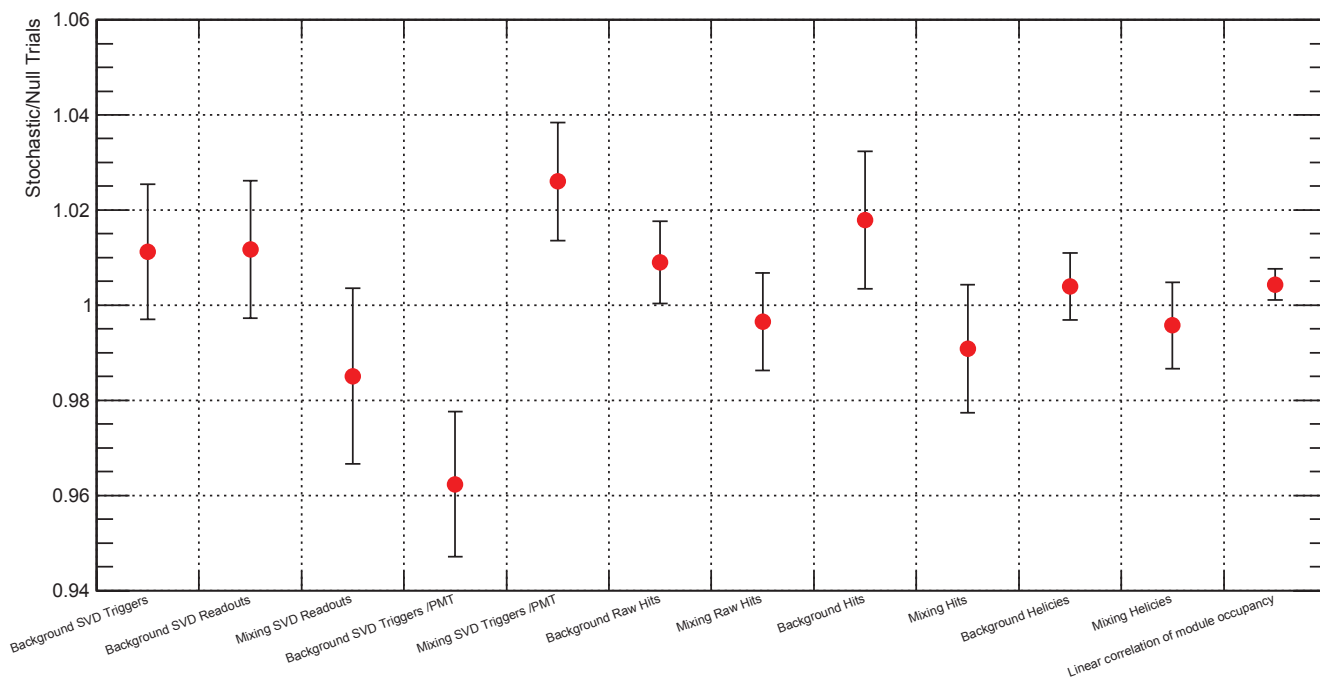
Extended Data Figure 1 | Quality-of-data (QOD) tests. QOD tests comparing the ratios of various quantities during stochastic and null trials. No difference corresponds to a ratio of 1, and the error bars are 1σ. Both mixing and background data were analysed in several tests. Background SVD triggers refers to the number of times the SVD meets our trigger condition. This number is counted even when the detector is vetoed because it is being read out. Background/mixing SVD readouts refers to the number of events read out by the detector. During cosmic (background) data collection, this value is not saturated; however, during mixing it is. Background/mixing SVD triggers/PMT is the ratio of the number of readout events that trigger the SVD, and the number of counts in adjacent scintillating detectors[3]. The scintillators have a lower detection efficiency and faster response than the SVD, so this ratio tests for saturation of the SVD rate. Background/mixing raw hits is a count of strip clusters with energy deposited 3σ above noise. This is the lowest level of reconstruction analysis. Background/mixing hits is a count of voxel hit clusters, that is, the number of coincident raw hits on both the n- and p-side of the silicon detector modules. Background/mixing helices is the number of three-hit combination tracks that point to the centre of the SVD. Linear correlation of module occupancy is the gradient of a line of best fit between the relative occupancy counts in the two experiments. Here occupancy is defined as the number of hits in a module.

Extended Data Table 1 | Results of MVA during the stochastic/hold phase (2)

Trial Type	Events
Stochastic Trials	6
Null Trials	11
Predicted Background	6.9

See Table 1 for the observed events during the release phase (3).