Journal of Biophysical Chemistry, 2026, 17(2), 15-27
https://www.scirp.org/journal/jbpc
ISSN Online: 2153-0378
ISSN Print: 2153-036X

Excess Correlation in pH between Magnetically Stimulated Beakers of Water: A Preregistered Replication

Dean Radin, Erik Brinsmead

Institute of Noetic Sciences, Novato, CA, USA
Email: dradin@noetic.org

How to cite this paper: Radin, D. and Brinsmead, E. (2026) Excess Correlation in pH between Magnetically Stimulated Beakers of Water: A Preregistered Replication. *Journal of Biophysical Chemistry*, **17**, 15-27.
https://doi.org/10.4236/jbpc.2026.172002

Received: April 18, 2026
Accepted: May 17, 2026
Published: May 20, 2026

Abstract

A preregistered replication was conducted to test a previously reported "excess correlation" effect using an automated microcontroller-based apparatus. Fifty experimental sessions were run where acetic acid was automatically added to a "local" beaker while magnetic stimulation was applied; these sessions were interleaved with 50 sham sessions using the same apparatus, but where no acid was added. The prediction was that pH in a "remote" beaker, where nothing was added, would shift toward alkaline during a specific phase of the magnetic stimulus. Analysis confirmed a statistically significant shift in pH in the remote beaker, as compared to a simultaneously measured control beaker, within two minutes of the predicted stimulus ($\Delta pH = +0.0034$, $p = 0.005$). By contrast, sham runs showed no effect ($\Delta pH = -0.0007$, $p = 0.53$). The observed pH shift was similar to a pilot study's finding ($\Delta pH = +0.004$), supporting the reproducibility of previously reported excess correlation effects in water pH and validating the use of an automated methodology for systematic investigation of this phenomenon.

Keywords

Excess Correlation, Magnetic Stimulation, pH, Water, Automated Experimentation, Sham Controls, Nonlocal Coupling

1. Introduction

The term "excess correlation" refers to an anomalous correlation said to arise between properties of physically isolated systems when both are exposed to the same sequence of magnetic field pulses [1]. Published studies have documented such effects across diverse physical targets ranging from pH in water, to seed germination, growth in cell cultures *in vitro*, photochemical reactions, and human brain-to-brain correlations [2]-[5].

These "entanglement-like" effects are regarded as controversial because the underlying mechanisms remain inadequately explained by conventional theories, and because they have been reported only by researchers from a single laboratory at Laurentian University. Despite multiple theoretical proposals, ranging from quantum entanglement to nonlocal informational fields [4] [6]-[8], there is no consensus on an adequate explanatory model.

A pilot replication of this phenomenon by the first author demonstrated a statistically significant pH shift in isolated water beakers exposed to the same weak (< 10 nT) magnetic stimuli ($\Delta pH = +0.004$, $p < 0.0005$) [9], consistent with previous reports from the Laurentian laboratory. The present study attempted to replicate that outcome by implementing an automated apparatus to eliminate potential experimenter influence, by having laboratory personnel other than the first author run the experiment in a distant laboratory, by preregistering the predicted outcome (osf.io/nt76h/overview), and by including a sham condition in which the identical equipment, protocol, and statistical methods were used except that no acid was added to the local beaker.

2. Methods

2.1. Equipment and Protocol

The apparatus consisted of two Arduino microcontroller boards, each equipped with SD card writers and battery-powered datetime clocks. Board A simultaneously energized the two toroids (called a "halo" because of the toroidal shape). Each halo consisted of a 25.4 cm diameter plastic hoop wound with 225 loops of 16-gauge copper wire. It generated ~20 nT magnetic fields with specific pulse sequences called *primer* and *effector*; pulse sequence details are described in [10]. Board B collected the pH sensor and temperature data and controlled the liquid handling pump (EZO-PMP, Atlas Scientific).

Three 40 mL glass beakers filled with tap water were labeled Local, Control, and Remote. This water, from Novato, California, may be characterized as a low-TDS (total dissolved solids), slightly alkaline (pH ~7.5 - 8.5), mildly hard (~82 mg/L $CaCO_3$) treated surface water from local reservoirs that has been chlorinated, pH-stabilized for distribution, and with low trace metals (e.g., arsenic in sub-ppb to low ppb range). It exhibits the expected trace disinfection byproducts and low-level naturally occurring metals typical of compliant U.S. municipal systems.

The Local beaker was placed in the center of one halo; it was the recipient of acid in experimental trials (acetic acid in white vinegar). The Remote beaker was placed in the center of the second halo, and nothing was added to it. The Control beaker was placed between the Local and Remote beakers as an environmental witness; it was not exposed to magnetic stimulation, nor was anything added to it. The testing apparatus is shown in **Figure 1**.

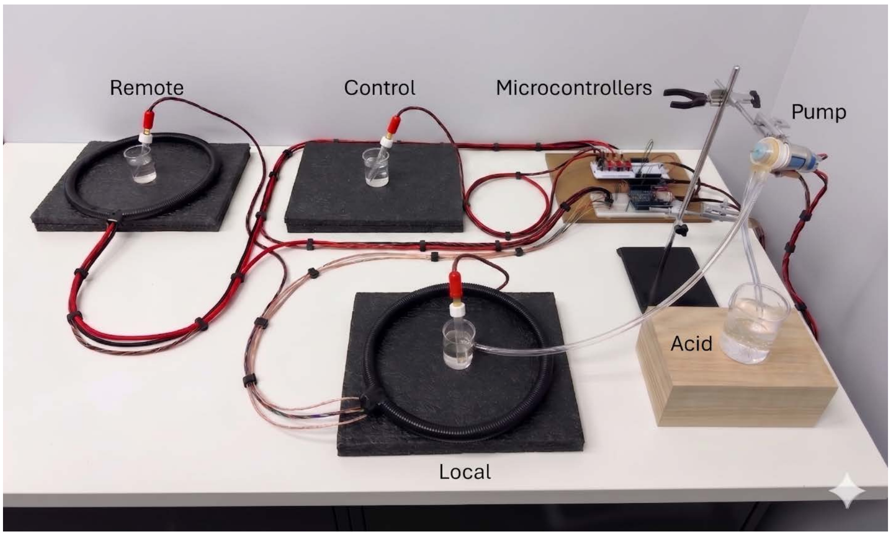

Figure 1. Layout of the automated testing system. One microcontroller energized the magnetic toroids around the Local and Remote beakers. A second microcontroller operated the pump that dropped acid into the Local beaker at the start of the effector magnetic stimulation sequence (see text). All three beakers simultaneously measured pH in water and the resulting data were stored at 0.5 Hz on an SD card. A fully automated run on this system was started by pressing a button on one of the microcontrollers. Fresh tap water was used to rinse and refill each beaker between sessions.

Water in each beaker was monitored by a pH sensor (Gen3, 0.001 pH resolution, Atlas Scientific, Long Island City, NY) at 0.5 Hz. Each sensor was calibrated using standard solutions (pH 4.0, 7.0, and 10.0), and according to the Atlas Scientific datasheet for these sensors (V5.1, revision 2/24), for weak levels of acids and bases recalibration was only necessary once per year over the first two years of usage (this study was completed within that timeframe).

A microcontroller operated a liquid dispensing pump (Atlas Scientific EZO-PMP), and a milligauss meter was used during development of the apparatus to measure the magnetic fields produced by the halos (AlphaLab Milligauss Meter, 1 nT resolution). Data from each run was automatically recorded on a 64 GB SD card.

A total of 100 runs were pre-planned, 50 each in interleaved experimental and sham conditions. The sham condition was conducted identically to the experimental condition except it deposited acid into a separate beaker located next to the Local beaker, providing roughly the same distance and orientation to the Remote beaker. This was designed to explicitly make the electromagnetic, thermal, and vibrational factors the same as in the experimental runs, except no acid was dropped into the Local beaker. A single run in this experiment was 30 minutes, consisting of an initial 6-minute period termed a *baseline* phase with no magnetic stimulus, followed by a 6-minute *primer* magnetic sequence, followed by a 12-minute *effector* magnetic sequence, and ending with a 6-minute baseline with no magnetic stimulus. All of these runs were conducted at the Institute of Noetic Sciences' laboratory by research assistants.

The Remote beaker included a temperature sensor (Atlas Scientific SMA PT-1000) that was used to automatically compensate for ambient temperature drift. The Control beaker's probe was not temperature-compensated, but the metric of interest was unaffected because the primary dependent variable was not raw pH, but rather a change in a difference metric. Specifically, $D(t)$ = pH_remote(t) − pH_control(t) was first computed at each time point of interest, and then $\Delta D = D\_post - D\_pre$ was derived as the change in that difference between the 2-minute post-event (*i.e.*, the acid drop) window and the 2-minute pre-event window. For asymmetric temperature compensation to produce a spurious ΔD, it would have required a sharp change in temperature that occurred precisely at the moment of acid delivery and that affected the two beakers differently. This possibility was physically implausible because the acid was only added to the Local beaker, and both the Remote and Control beakers were in the same ambient environment, so they experienced the same temperature trajectory. Because any temperature-driven drift in the uncompensated Control probe was common to both the pre-event and post-event windows, it was subtracted out in ΔD. Slow monotonic drift, which would have been the dominant form of ambient temperature changes, would leave ΔD unchanged by construction, which was the virtue of using a difference-of-differences statistic.

2.2. Hypothesis

The excess correlation hypothesis predicted that when acetic acid was added to the Local beaker of water, the pH of water in the Remote beaker would shift toward alkaline when both toroids were exposed to the same magnetic stimulation sequence. Specifically, the hypothesis predicted a magnetic phase-specific and temporally-localized effect, which had been observed in the pilot study and demonstrated as a time-window-specific scalar analysis, as noted in that publication [9]. The preregistered analysis thus focused on a change in pH in the Remote beaker during the transition from the *primer* to the *effector* phase.

2.3. Analyses

2.3.1. Session Inclusion

All recorded sessions were included, provided that valid pH data were recorded in the Local, Remote, and Control beakers, that time stamps were present and monotonically increasing, and that all data samples were available within the session and recorded properly on the SD card. Session data were not excluded based on Remote or Control outcomes.

2.3.2. Automated Classification of Experimental vs Sham Runs

Sessions were classified automatically as experimental or sham by the analytical code by measuring the Local beaker's change in pH when the pump was programmed to dispense acid. To accomplish this, the median pH in the final 60 s of the session was subtracted from the median pH in the first 60 s. If the resulting difference was ≤ -2.0 pH units, the session was classified as experimental, otherwise it was classified as sham. This criterion corresponded to the known shift in pH due to acid delivery and did not reference pH in the Remote or Control beaker data.

2.3.3. Event Time Detection

For experimental sessions, acid delivery was programmed to occur at 720 seconds after the start of a run, which was when the *effector* magnetic stimulus phase began. The actual acid drop event time was defined as the moment when the largest negative step was observed in a smoothed local pH signal. That detected time was used as the "event anchor," which is described in more detail below. It typically took about 5 to 10 seconds after the initial acid drop for pH to reach a maximum.

The drop-time detection procedure operated as follows. The local pH signal was first smoothed with a 3-point moving median (window = 6 s at the 0.5 Hz sampling rate) applied within a [600, 900] s search window. The first difference of the smoothed signal was then computed, and the event anchor was defined as the time point corresponding to the single largest negative step in that differenced series, *i.e.*, the steepest downward transition in local pH within the search window. No threshold criterion was applied; the method simply identifies the most abrupt drop.

For the primary analysis, a "hybrid" anchor mode was used: experimental runs used the detected drop time, while sham runs used a fixed anchor of 730 s. The 730 s value was chosen rather than the programmed 720 s to account for the observed ~10 s lag between pump activation and the onset of a measurable pH step, as noted in Section 2.3.3.

As a robustness check, the primary analysis was repeated using a fixed anchor of 720 s for all runs, bypassing the event-detection step entirely. The result was $\Delta D = +0.0041$, $p = 0.005$, indistinguishable from the primary result ($\Delta D = +0.0041$, $p = 0.005$). This was expected given that the median detected drop time across experimental runs was 730.0 s, meaning the two anchor choices differed by only 10 seconds relative to a ±120 s analysis window. The primary finding therefore did not depend on the event-detection procedure.

2.3.4. Primary Dependent Variable

The dependent variable was the difference in pH in the Remote vs. Control beaker

$$D(t) = \mathrm{pH}_{\mathrm{remote}}(t) - \mathrm{pH}_{\mathrm{control}}(t)$$

This difference was intended to reduce shared drift, ambient temperature effects, and sensor instabilities. The Control refers to the beaker that was neither magnetically stimulated nor had acid added.

2.3.5. Pre-Specified Confirmatory Analysis

A single scalar value was computed for each session as a differential measure

$$\Delta D = \bar{D}_{\mathrm{post}} - \bar{D}_{\mathrm{pre}}$$

within a fixed window of time. The post-event window was defined as 0 to 2 minutes relative to the event anchor, and the pre-event window as the 2 minutes prior to the event anchor. The anchoring scheme for the experimental runs was based on the moments when acid was delivered, and for sham runs it was anchored to the same time. This anchoring scheme, which determined when the predicted pH change would occur, was fixed *a priori*. The ΔD value was then tested against chance outcome via permutation analysis (20,000 permutations), yielding a one-tailed *p*-value.

As a secondary analysis, we also fit a linear mixed-effects model to see if the effect generalized to broader epoch structures: ΔpH ~ Time_c + Epoch*Acid + (1 + Time_c|RunID). This model tested if the shift in pH extended across the entire 6-minute *primer* and 12-minute *effector* phases.

2.3.6. Sensitivity Grid with Family-Wise Error-Rate Correction

The primary scalar ΔD depended on two choices by the analyst: the boundaries of the pre-event window and the boundaries of the post-event window, each specified in seconds relative to the anchor time. To examine the role of these choices, a sensitivity grid was included in the analysis plan to address three distinct concerns that a single preregistered window choice could not address on its own. First, it provided robustness information, such that if the primary result held only for the exact preregistered window but collapsed for neighboring windows, then that pattern would suggest the effect was not robust and would warrant a different interpretation than one in which the result was stable across a broader region of the parameter space. Second, it provided descriptive information about the temporal shape of any effect that was present, which a single window could not resolve. Third, it provided a principled multiplicity-adjusted *p*-value that a reader could substitute for the preregistered *p*-value if they preferred to treat the window choice as exploratory rather than fixed in advance (even though it was not exploratory). The grid was thus not a confirmatory test, but rather it supplemented it as a way to explore the investigators' analytical flexibility.

Design

The grid was formed by crossing a set of candidate pre-windows with a set of candidate post-windows. Pre-window onsets ranged from −240 s to −60 s relative to the anchor, with widths from 60 s to 220 s. Post-window onsets ranged from 0 s to +30 s, with widths from 60 s to 210 s. This yielded 81 window pairs spanning a plausible neighborhood of the preregistered choice for where the primary effect of interest would appear. For each pair, the scalar ΔD was recomputed for every run using that pair's pre- and post-window boundaries, and the between-group mean difference (experimental minus sham) was recorded along with a one-tailed nominal permutation *p*-value. These *p*-values were not themselves interpretable as evidence, because 81 tests were performed simultaneously and any one of them might have attained a small *p*-value by chance.

To obtain valid inference across the grid, a max-statistic permutation procedure was used. On each of 2000 permutations the condition labels were shuffled across runs, and under the shuffled labels the between-group mean difference was recomputed for each of the 81 window pairs. From each permutation, only the single largest of those 81 differences was retained, and this yielded a null distribution of grid-wide maxima. For any given window, the family-wise-error-rate or FWER-corrected p-value was then the proportion of permutations in which this grid-wide maximum equaled or exceeded the observed between-group difference at that window. This procedure controlled the family-wise error rate at the nominal α level, meaning the probability of declaring any window significant under the null hypothesis remained bounded by α, regardless of how many windows were examined. It also accommodated the correlation structure among adjacent windows automatically, without requiring that structure to be explicitly modeled.

Two p-values were reported for each window: the nominal p-value, which treated that window as if it had been analyzed in isolation, and the FWER-corrected p-value, which treated the window as one of a family of 81 post-hoc selections from the grid. In addition, a single global FWER-corrected p-value was reported, equal to the corrected p-value of the window with the smallest nominal p-value. This answered the question, "was there *any* window in the grid at which the experimental group reliably exceeded the sham group after correcting for the search."

Interpretive framework

Several features of the resulting grid output were informative. A global FWER-corrected p-value below α would indicate that at least one window in the grid showed a reliable effect after adjustment for performing multiple tests. The count of windows surviving FWER correction, and the direction of their effects, would indicate if any reliable signal was coherent or scattered. A genuine time-locked effect was expected to produce a contiguous region of windows whose observed differences all pointed in the predicted direction, whereas noise was expected to produce an approximately symmetric scatter of nominally significant windows in both directions at approximately the nominal rate. The location in parameter space where large differences concentrated, if any such cases existed, was descriptive of the temporal shape of the effect, e.g., was it too narrow and tightly centered on the anchor, or broader and sustained, and could it be compared to the preregistered window to judge if the preregistration captured the effect near its optimum. Finally, the contrast between the preregistered window's nominal p-value and its FWER-corrected p-value quantified the penalty that post hoc selections of that same window would have incurred. The preregistered window retained its status as the primary confirmatory inference regardless of how the grid behaved. The grid results are reported to allow readers to form their own judgments about the robustness, shape, and multiplicity-adjusted evidential weight of the primary finding.

3. Results

3.1. Primary Result

The preregistered primary outcome was a scalar, ΔD, built from the Remote-Control pH difference in a post-event window (0 to +120 s around the moment of acid delivery) minus its mean in a pre-event window (−120 to 0 s before acid delivery). Inference was a one-tailed permutation test (20,000 permutations) comparing group means, with the directional prediction that the experimental group would exceed the sham group.

Experimental runs showed a mean ΔD of +0.0034 pH units; sham runs showed −0.0007. The between-group difference was +0.0041, and the one-tailed permutation result was $p = 0.005$. This meant that after acid was delivered to the local beaker, the Remote beaker significantly drifted toward alkaline relative to the unstimulated control, as predicted. This small magnitude but reliable shift was absent in sham runs (see **Figure 2**). The direction and magnitude of the change in pH both matched the results observed in the previously published experiment [9].

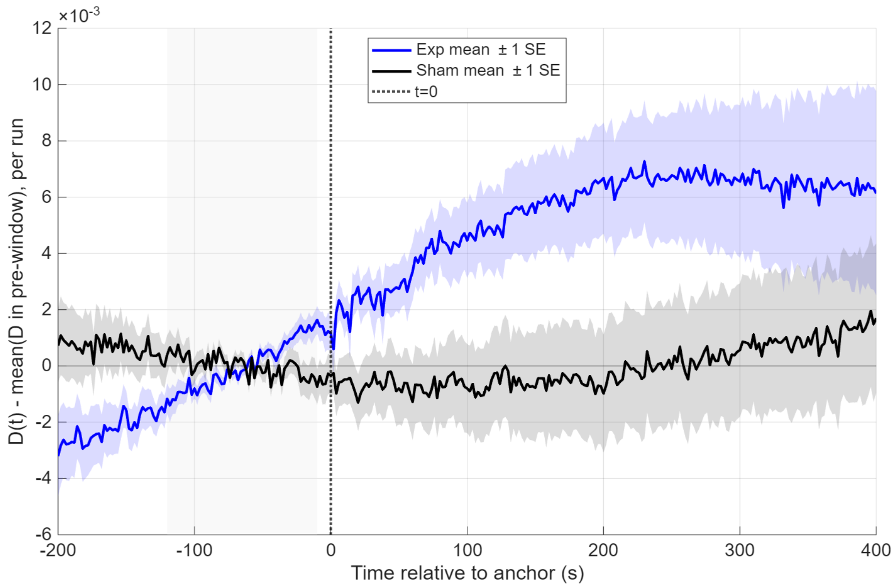

Figure 2. Change in pH in experimental runs and sham runs, with one standard error bars, from 200 seconds before to 400 seconds after the acid drop (at time 0). The preregistered prediction was a rise in pH in the Remote water after the *effector* sequence began, and no change in the pH of the Control water.

Two additional analyses were designed to investigate possible spurious signals due to general pH drift. Both produced null results when re-applying the same primary test but shifted in time when the analysis windows were placed. Recall that the "anchor" was the point in absolute session time that was treated as the event of interest; the scalar ΔD was defined as the mean of $D(t)$ in the 120 s *after* the anchor minus the mean of $D(t)$ in the 120 s *before* it. In the primary analysis, the anchor was the detected acid-delivery time (~730 s into each recording, shown as t = 0 in **Figure 2**).

In one far-anchor negative control (meaning no effect should appear), the anchor point was moved to 300 s of absolute recording time, roughly 430 s *before* the actual acid drop, well to the left of the range plotted in **Figure 2**. The ±120 s windows around it, and each run's Experimental/Sham classification, were left unchanged. This procedure tested if a differential effect would appear if the event happened 430 s earlier than the actual anchor point, and it did not ($p = 0.26$).

In a far-pre window negative control, the real acid-drop anchor was retained, but both the "pre" and "post" windows were slid before the event (to [−600, −480] s and [−480, −360] s relative to the anchor). This investigated a complementary question, namely was there a comparable Experiment-vs-Sham difference in the quiet interval just preceding acid delivery? Again, the answer was no ($p = 0.27$).

Taken together, these tests confirmed that the effect was localized to the acid-delivery moment itself, whereas arbitrary slices of the session, whether taken hundreds of seconds earlier in absolute time, or in the minutes immediately preceding the drop, did not reproduce it. Only the windows flanking the actual acid delivery time yielded the significant Experimental-vs-Sham separation reported for the primary outcome.

3.2. Secondary Test

The secondary test was a linear mixed-effects regression of the Remote vs. Control difference on the factors of Time, Epoch (primer vs. effector), Condition, and their interaction, with per-session random intercepts and slopes. The sample-level fit had autocorrelated residuals (lag-1 ACF ≈ 0.97 at 2-second sampling), which inflated apparent precision, so two autocorrelation-robust variants were conducted. One was a run-level aggregated LME (one observation per run per epoch) and the second was a cluster bootstrap which resampled whole runs. Both results showed the same null effect, indicating that the Epoch × Condition interaction was not significant ($p \approx 0.48$, one-tailed).

This means the effect was not sustained during the 12-minute *effector* epoch. It was instead a narrow, transient excursion concentrated in a ±120-second window around acid delivery, which is what the primary scalar was designed to detect and which was averaged away when the LME integrated over the whole epoch.

3.3. Sensitivity Grid

The formal inference was anchored to the preregistered window around the anchor point of interest, but a grid of 81 pre/post window combinations was also evaluated with a max-statistic permutation test (2000 permutations) to control the family-wise error rate (FWER), and to test how robust the predicted results were. The results showed that the global FWER $p = 0.018$, which was corrected for multiple comparisons after testing the 81 different window choices. Fifteen of the 81 windows survived FWER correction at $p < 0.05$, and all of them showed a positive difference (experimental > sham). No window showed a significant effect in the opposite direction, indicating a genuine underlying signal rather than a symmetric distribution of noise.

3.4. Session Exclusions and Condition Sequencing

One session was excluded because not all of the pH samples were properly recorded, and in three sham sessions acid was accidentally dropped into the Local beaker. If a session was excluded another session was run to ensure that the final analysis included exactly 50 experimental and 50 sham runs. The sequence of session types was intended to alternate between sham and experimental, but this was not strictly followed in the first 35 runs due to a miscommunication with the laboratory assistant who ran the sessions. Starting with session 36, all conditions were properly alternated. To then check if the sequence of session types in the first 35 runs (call this phase 1) might have introduced a monotonic drift artifact, we compared the results of those runs with the last 65 runs (phase 2).

That analysis showed that phase 1 resulted in ΔpH = +0.00496, $p = 0.063$, and phase 2 resulted in ΔpH = +0.00359, $p = 0.023$. If the primary outcome of $p = 0.0052$ were an artifact of slow drift, we would have expected the phase 1 outcome to show strong significance and phase 2 to be null. Instead, both phases showed positive pH differences of similar magnitude, thus, there was no indication that the phase 1 sessions were due to the presence of a simple monotonic drift.

3.5. Run-Order Drift

To address the possibility of run-order or batch-level drift, we conducted two analyses. First, Spearman correlations between chronological run index and the scalar ΔD were computed separately within each condition. Neither the experimental runs ($\rho = -0.056$, $p = 0.70$) nor the sham runs ($\rho = -0.142$, $p = 0.33$) showed a meaningful association with run order. Second, run index was added as a continuous covariate in a linear regression predicting ΔD. The run-order coefficient was negligible and non-significant ($\beta = -0.000026$, $p = 0.34$), while the acid effect remained essentially unchanged ($\beta = +0.0041$, $p = 0.012$), identical to the unadjusted estimate. Together with the phase-split analysis already reported in Section 3.4, which showed similar effect magnitudes in both the early blocked phase ($\Delta D = +0.0050$) and the later alternating phase ($\Delta D = +0.0036$), these results provide no evidence that the primary finding was attributable to run-order drift or batch effects.

4. Discussion

The previously published pilot study reported the same shift in pH (toward alkaline) and a need for autocorrelation-aware statistics [9]. The present automated analysis achieved a similar conclusion via a different statistical route, and a preregistered permutation test on a run-level scalar, with $p = 0.005$. What this means is that across 100 preregistered runs, with automated classification, negative controls, an autocorrelation-robust inference, and conducted in a different laboratory than the pilot study with different personnel, there was a statistically reliable, time-locked, small-amplitude alkaline shift in the remote beaker's pH at the moment acid was added to the physically isolated local beaker. This shift was absent in sham runs and absent at other anchor times.

Regarding pH sensor stability, the Atlas Scientific Gen3 probes used in this study were Ion-Sensitive Field-Effect Transistor (ISFET)-based rather than glass electrodes. ISFET sensors under low ionic strength conditions, such as those in the present experiment, have substantially lower drift than glass electrodes [11], and studies of glass probes have shown that drifts accumulate over weeks to months, not within the 30-minute sessions used here [12]. Together, this supports the adequacy of the pH sensors used in the present study to detect small changes in pH, provided that a difference-of-differences analysis structure is used (as it was).

Another possible concern is the high temporal autocorrelation typical of slow electrochemical processes, which renders naive repeated-measures approaches invalid. The run-level scalar approach used as the primary inference strategy in this study, which collapsed the full time series to a single number per session before group comparison, is a standard solution to this problem [13]. Complementary validation using cluster bootstrap resampling of whole runs further guarded against residual within-session dependence.

Regarding replication of anomalous claims like "excess correlation," methodological recommendations including preregistration, automation to reduce experimenter influence, sham controls, independent site replication, and transparent reporting, have all been articulated in recent literature [14] [15], and the present study was designed to meet such requirements.

What remains missing is an adequate explanatory mechanism. The magnetic field involved was ~10 nT, the beakers were about a meter apart, and no conventional electromagnetic, thermal, or chemical coupling route would have been the obvious cause for the observed effect. Thus, while this result replicates a previously published anomalous finding using a more robust method, independent replication by other laboratories is essential to more convincingly establish the existence of a genuine anomaly.

4.1. Limitations

One limitation in this study was that the magnitude of the observed effect was near the pH sensor's stated resolution (0.001 pH). While the observed effect was statistically above this resolution, caution is warranted.

Another limitation was the temporal autocorrelation in the pH signal. While not unexpected for continuous, slow-moving chemical processes, this precluded use of simple statistical methods. Fortunately, the scalar approach used in the present replication captured the effect in its natural temporal window, which partially mitigates this limitation.

A third limitation was that the sensitivity grid test, which was designed to verify specificity of the predicted effect, necessarily included a large set of possibilities. While the FWER approach provided a viable way to correct for multiple tests, future studies might employ more restrictive grids to further reduce concerns about multiple comparisons.

4.2. Future Tests

If the excess correlation effect in pH is indeed genuine, then mechanistic investigations become a priority. In that regard, besides encouraging independent replication, future studies should a) vary the magnetic stimulation parameters (frequency, intensity, waveform) to identify optimization conditions; b) more precisely characterize the water properties (mineral content, initial pH, temperature); c) systematically examine distance limits of the effect; d) investigate alternative explanations (electromagnetic contamination, air circulation, temperature coupling) with environmental witness sensors; e) measure additional variables (conductivity, dissolved oxygen, temperature) to assess specificity to pH, and f) develop testable theoretical models compatible with these observations. Some efforts along these lines were reported by members of Persinger's laboratory (e.g. [1] [10]), but again, truly independent efforts are indispensable.

5. Conclusion

This study replicated a small but statistically significant excess correlation effect in pH between isolated beakers of water exposed to the same pulsed magnetic stimulation sequences using an automated apparatus located in a laboratory operated by personnel not previously engaged in this line of research. The phenomenon remains mechanistically unexplained and warrants further investigation by the broader scientific community.

Acknowledgements

This work was supported by the Bial Foundation. The authors acknowledge assistance by Richard Tian in running portions of this experiment, and the conceptual foundations from Michael Persinger's laboratory. MATLAB analysis scripts and raw data are available from the author upon request.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

[1] Rouleau, N., Carniello, T.N. and Persinger, M.A. (2016) Identifying Factors Which Contribute to the Magnitude of Excess Correlations between Magnetic Field-Paired Volumes of Water. *Journal of Signal and Information Processing*, **7**, 136-147. https://doi.org/10.4236/jsip.2016.73014

[2] Burke, R.C., Gauthier, M.Y., Rouleau, N. and Persinger, M.A. (2013) Experimental Demonstration of Potential Entanglement of Brain Activity Over 300 Km for Pairs of Subjects Sharing the Same Circular Rotating, Angular Accelerating Magnetic Fields: Verification by s_LORETA, QEEG Measurements. *Journal of Consciousness Exploration & Research*, **4**, 35-44.

[3] Rouleau, N., Carniello, T.N. and Persinger, M.A. (2014) Non-Local pH Shifts and Shared Changing Angular Velocity Magnetic Fields: Discrete Energies and the Importance of Point Durations. *Journal of Biophysical Chemistry*, **5**, 44-53. https://doi.org/10.4236/jbpc.2014.52006

[4] Dotta, B.T., Mulligan, B.P., Hunter, M.D. and Persinger, M.A. (2009) Evidence of

Macroscopic Quantum Entanglement during Double Quantitative Electroencephalographic Measurements of Friends vs Strangers. *NeuroQuantology*, **7**, 548-551. https://doi.org/10.14704/nq.2009.7.4.251

[5] Dotta, B.T. and Persinger, M.A. (2012) "Doubling" of Local Photon Emissions When Two Simultaneous, Spatially-Separated, Chemiluminescent Reactions Share the Same Magnetic Field Configurations. *Journal of Biophysical Chemistry*, **3**, 72-80. https://doi.org/10.4236/jbpc.2012.31009

[6] Karbowski, L.M., Murugan, N.J. and Persinger, M.A. (2015) Experimentally-Induced Inhibition of Growth in Melanoma Cell Cultures Separated by ~2 Kilometers When Both Share Excess Correlation Magnetic Fields: Macroscopic Evidence of Free-Space Quantum Teleportation? *Journal of Signal and Information Processing*, **6**, 39-48. https://doi.org/10.4236/jsip.2015.61004

[7] Persinger, M.A., Saroka, K.S., Lavallee, C.F., Booth, J.N., Hunter, M.D., Mulligan, B.P., *et al.* (2010) Correlated Cerebral Events between Physically and Sensory Isolated Pairs of Subjects Exposed to Yoked Circumcerebral Magnetic Fields. *Neuroscience Letters*, **486**, 231-234. https://doi.org/10.1016/j.neulet.2010.09.060

[8] Persinger, M.A. and Koren, S.A. (2007) A Theory of Neurophysics and Quantum Neuroscience: Implications for Brain Function and the Limits of Consciousness. *International Journal of Neuroscience*, **117**, 157-175. https://doi.org/10.1080/00207450500535784

[9] Radin, D. (2025) Independent Replication of an "Excess Correlation" Effect in pH between Isolated Beakers of Water. *Journal of Biophysical Chemistry*, **16**, 15-29.

[10] Rouleau, N. and Persinger, M. (2015) Local Electromagnetic Fields Exhibit Temporally Non-Linear, East-West Oriented 1 - 5 NT Diminishments within a Toroid: Empirical Measurement and Quantitative Solutions Indicating a Potential Mechanism for Excess Correlation. *Journal of Electromagnetic Analysis and Applications*, **7**, 19-30. https://doi.org/10.4236/jemaa.2015.72003

[11] Bagshaw, E.A., Wadham, J.L., Tranter, M., Beaton, A.D., Hawkings, J.R., Lamarche-Gagnon, G., *et al.* (2021) Measuring pH in Low Ionic Strength Glacial Meltwaters Using Ion Selective Field Effect Transistor (ISFET) Technology. *Limnology and Oceanography: Methods*, **19**, 222-233. https://doi.org/10.1002/lom3.10416

[12] Wiora, A. and Wiora, J. (2018) Over One-Year Long-Term Laboratory Tests of Ph Electrodes in Terms of Industrial Applications Checking Stabilities of Their Parameters and Their Influence on Uncertainties of Measurements. *Sensors*, **18**, Article 4102. https://doi.org/10.3390/s18124102

[13] Hounyo, U. and Kim, M.S. (2025) Robust Two-Sample Mean Inference under Serial Dependence. arXiv:2512.11259.

[14] Reber, A.S. and Alcock, J.E. (2020) Searching for the Impossible: Parapsychology's Elusive Quest. *American Psychologist*, **75**, 391-399. https://doi.org/10.1037/amp0000486

[15] Waldron, S. and Allen, C. (2022) Not All Pre-Registrations Are Equal. *Neuropsychopharmacology*, **47**, 2181-2183. https://doi.org/10.1038/s41386-022-01418-x