

The perception of silence

Rui Zhe Goh, Ian B. Phillips^{*†}, Chaz Firestone^{*†}

Affiliations:

Department of Psychological and Brain Sciences, Johns Hopkins University; Baltimore, Maryland, 21218, USA

Department of Philosophy, Johns Hopkins University; Baltimore, Maryland, 21218, USA

†Phillips and Firestone contributed equally and should be considered joint senior author.

*Corresponding authors: Ian Phillips and Chaz Firestone

Email: ianbphillips@jhu.edu, chaz@jhu.edu

Version: 5/10/23; in press, *Proceedings of the National Academy of Sciences*

Abstract

Auditory perception is traditionally conceived as the perception of *sounds* — a friend's voice, a clap of thunder, a minor chord. However, daily life also seems to present us with experiences characterized by the *absence* of sound — a moment of silence, a gap between thunderclaps, the hush after a musical performance. In these cases, do we positively *hear silence*? Or do we just *fail to hear*, and merely judge or infer that it is silent? This longstanding question remains controversial in both the philosophy and science of perception, with prominent theories holding that sounds are the only objects of auditory experience and thus that our encounter with silence is cognitive, not perceptual. However, this debate has largely remained theoretical, without a key empirical test. Here, we introduce a novel empirical approach, presenting experimental evidence that silence can be genuinely perceived (not just cognitively inferred). We ask whether silences can 'substitute' for sounds in event-based auditory illusions — empirical signatures of auditory event representation in which auditory events distort perceived duration. Seven experiments introduce three new 'silence illusions' — the one-silence-is-more illusion, silence-based warping, and the oddball-silence illusion — each adapted from a prominent perceptual illusion previously thought to arise only from sounds. Subjects were immersed in ambient noise interrupted by silences structurally identical to the sounds in the original illusions. In all cases, silences elicited temporal distortions perfectly analogous to the illusions produced by sounds. Our results suggest that silence is truly heard, not merely inferred, introducing a general approach for studying absence perception.

Significance Statement

Do we only hear sounds? Or can we also hear silence? These questions are the subject of a centuries-old philosophical debate between two camps: The perceptual view (we literally hear silence), and the cognitive view (we only judge or infer silence). Here, we take a novel empirical approach to resolve this theoretical controversy. We show that silences can 'substitute' for sounds in event-based auditory illusions. Seven experiments introduce three new 'silence illusions', adapted from perceptual illusions previously thought to arise only with sounds. In all cases, silences elicited temporal distortions perfectly analogous to their sound-based counterparts, suggesting that auditory processing treats moments of silence the way it treats sounds. Silence is truly perceived, not merely inferred.

Main Text

Introduction

What do we hear? The canonical answer is that auditory perception is the perception of *sounds* and their properties — the pitch of a friend’s voice, the loudness of a thunderclap, the timbre of a minor chord. This traditional view has considerable pedigree, with influential historical sources holding that sounds are the sole objects of auditory perception (1; cf.2). It is also the answer favored in contemporary scholarship: Prominent scientific accounts conceive the fundamental units of auditory perception as sounds (or auditory streams comprised of sounds; 3,4), and many philosophical theories agree, holding that “all auditory perception involves the perception of sound” (5) and that “if anything at all is heard, what is heard is necessarily a sound” (6) (see also 7,8). The pervasiveness of this canonical view about the contents of auditory perception might seem unsurprising — what else might we hear, if not sound?

However, there has long been a stubborn and intuitive counterexample: experiences of *silence*, which are characterized by the absence of sound. Silence confronts us throughout our daily lives — consider an awkward pause in a conversation, a suspenseful gap between thunderclaps, or the hush at the end of a musical performance. What is the nature of these experiences?

Silence: Heard or inferred?

One possibility is that experiences of silence are simply cases in which we *fail to hear*, and then use our faculties of reasoning and judgment to *infer* that it is silent. This interpretation is offered by those who defend the traditional sound-only view of audition, holding that an experience of silence is merely the “cognitive accompaniment of an absence of experience” and “is itself no form of hearing” (9). This cognitive view may be motivated by a deeper assumption about perception, namely that we can genuinely perceive only what is *present* in the world, not what is absent (9,10). After all, one might think, absences are non-entities — they do not exist — and so can hardly impinge on our sensory apparatus.

However, an alternative possibility, which arguably does more justice to our phenomenology, is that we literally *perceive* silences. This interpretation has recently received support from philosophers who insist that hearing silence is not a mere failure to hear sound but rather a case of successful perception (11,12), and that despite their apparent attractions cognitive views are ill-motivated (13).

There is of course a rich empirical literature demonstrating the importance of pauses and gaps in auditory perception (e.g., speech processing (14-17) and word segmentation (18)); however, while these findings certainly enrich our understanding of the role that silence plays in the individuation and identification of *sounds*, they are neutral on whether *silences themselves* can be objects of auditory perception. The same is true of work on musical endings (19), vocal hesitation (20), and other phenomena, which investigate representations of sound-sequence

boundaries or breaks in speech, but not moments of silence themselves. Similarly, neuroscientific studies on sound termination (21), gap detection (22), and expectation violation (23) show that the brain is sensitive to sound offsets and omissions, but leave open whether these phenomena amount to the genuine perception of silence. As a result, extant empirical work does not resolve the broader theoretical debate concerning perceptual versus cognitive accounts of silence, leaving both sides to rely predominantly on introspection, thought experiments and philosophical theorizing (9,11,12,13).

An empirical approach: Substituting silences for sounds

Here, we introduce a new empirical approach aimed directly at this foundational question, by asking whether silences can ‘substitute’ for sounds in event-based auditory illusions. Our approach focuses on a key aspect of auditory processing: *event segmentation* (3,24). When presented with acoustic stimuli, the auditory system segments continuous input into discrete event representations which typically correspond to the sounds we hear (e.g., musical notes, phones, and other units of auditory processing). An empirical signature of this segmentation process is that auditory event representations can cause illusions in which perceived duration is distorted. We reasoned that if the auditory system treats silences as genuine auditory objects and constructs event representations on their basis, then periods of silence should elicit temporal distortions analogous to those elicited by sounds. In the seven experiments that follow, we report three novel ‘silence illusions’ where this is indeed the case — i.e., in which previously discovered auditory illusions occur even when the sounds are replaced by silences (Fig. 1A). These illusions not only generate robust empirical data, but can easily be subjectively experienced (see ‘demos’ at <https://perceptionresearch.org/silence>). We conclude that the auditory system generates perceptual event representations corresponding to moments of silence (henceforth, *representations of silence*), and that these representations underlie genuine perceptual experiences of silence.

Experiments 1–3: The ‘one-silence-is-more’ illusion

The first illusion we introduce is the *one-silence-is-more illusion* (Experiments 1–3). This illusion was adapted from the (sound-based) one-is-more illusion, in which a single continuous tone is perceived as longer than two discrete tones having the same total duration (25). The one-is-more illusion has its basis in event representation — one auditory event seems longer than an objectively equated sequence comprising two auditory events. Experiment 1 inverted this illusion by substituting silences for sounds: rather than ask whether one long tone is perceived as longer than a sequence comprising two short tones, we asked whether one long silence is perceived as longer than a sequence comprising two short silences. To create convincing periods of silence, we immersed subjects in realistic ambient noise (e.g., a busy restaurant, a loud playground, a bustling market) for several minutes during an instruction and practice phase (Fig. 1B and 1C). After immersion, periods of silence were introduced by briefly cutting out the ambient noise (for 1-5s). On each trial, subjects experienced two sequences of silence that were structurally identical to the sequences of sound in the original illusion: a *one-silence sequence* in which the ambient noise went silent *once*, and a *two-silences sequence* in which

the ambient noise went silent *twice*, resuming briefly between silences. Subjects judged which sequence was longer. On critical trials in which the two sequences were of equal objective duration, subjects consistently judged the one-silence sequence as longer than the two-silences sequences ($t(84)=6.96$, $p<.001$; Fig. 2A); moreover, the mean proportion of ‘one-silence longer’ responses was strikingly similar to ‘one-tone longer’ responses in the original one-is-more illusion (0.66 in both cases; 25), suggesting that the one-is-more illusion occurs with silences in just the way it does with sounds. Importantly, this result held even for trials in which both sequences had equal durations of silence (such that the two-silences sequence was objectively longer from start to finish; $t(78)=2.46$, $p=.016$), thus ruling out the potential confound that subjects might have been comparing the total time spent in silence instead of the total durations of the entire sequences. These initial results exemplify the research strategy adopted here: An illusion previously thought to depend on sounds also occurs when the sounds are replaced by silences, consistent with the hypothesis that silences can be represented as genuine auditory events.

Experiment 2 replicated the one-silence-is-more illusion with a different behavioral measure. This time, each trial contained only one sequence; instead of comparing sequence durations by answering which sequence lasted longer, subjects estimated the duration of the single sequence they had just heard by holding down a key to reproduce the duration. Reproduced durations were significantly longer for one-silence trials than two-silences trials ($t(96)=4.73$, $p<.001$), suggesting that our results are not an artifact of any one measure (Fig. 2B).¹

One might worry that the one-silence-is-more illusion is not caused by auditory event representations of silence, but is instead due to attentional effects such as distraction by the intervening noise. Experiment 3 addressed this confound by adding a third sequence type: the ‘occluded-silence’ sequence. Occluded silences contained an intervening noise that was not a resumption of ambient noise, but rather a different noise altogether (a bird chirping) which sounds like it is being played *over* a single continuous silence. The bird chirping noise and the resumption of ambient noise were equated for perceived loudness and duration; thus, any attentional effects of the intervening noise in the two-silences condition should also be present in the occluded-silence condition. Even when intervening noises were equated in this way, reproduced durations were still longer for occluded-silence trials than two-silences trials ($t(85)=2.92$, $p<.005$), demonstrating that the one-silence-is-more illusion is driven by a difference in number of perceptual event representations, not merely by the attentional effects of an intervening noise (Fig. 2C). Beyond ruling out such confounds, Experiment 3 also shows that representations of silence can persist through occlusion — a key feature of perceptual object representations (26–29). Just as we can see a single object persisting behind an occluder, we can hear a period of silence persisting through an occluding noise.

¹ To further ensure that the effects reported here do not reflect misunderstanding of the task (e.g., reproducing the total time spent in silence, or just one of the silences in the two-silences condition, instead of the total durations of the entire sequences), we also ran a follow-up experiment that was identical to Experiment 2 except that, at the conclusion of the experimental session, all subjects answered a debriefing question asking them to verify the task instructions using values from a sample trial. Even excluding any subject who failed to answer this question correctly (i.e., considering only those subjects who correctly verified the instructions), the one-silence-is-more illusion emerged, $t(80)=4.05$, $p<.001$. We thank a reviewer for comments that led to this follow-up experiment.

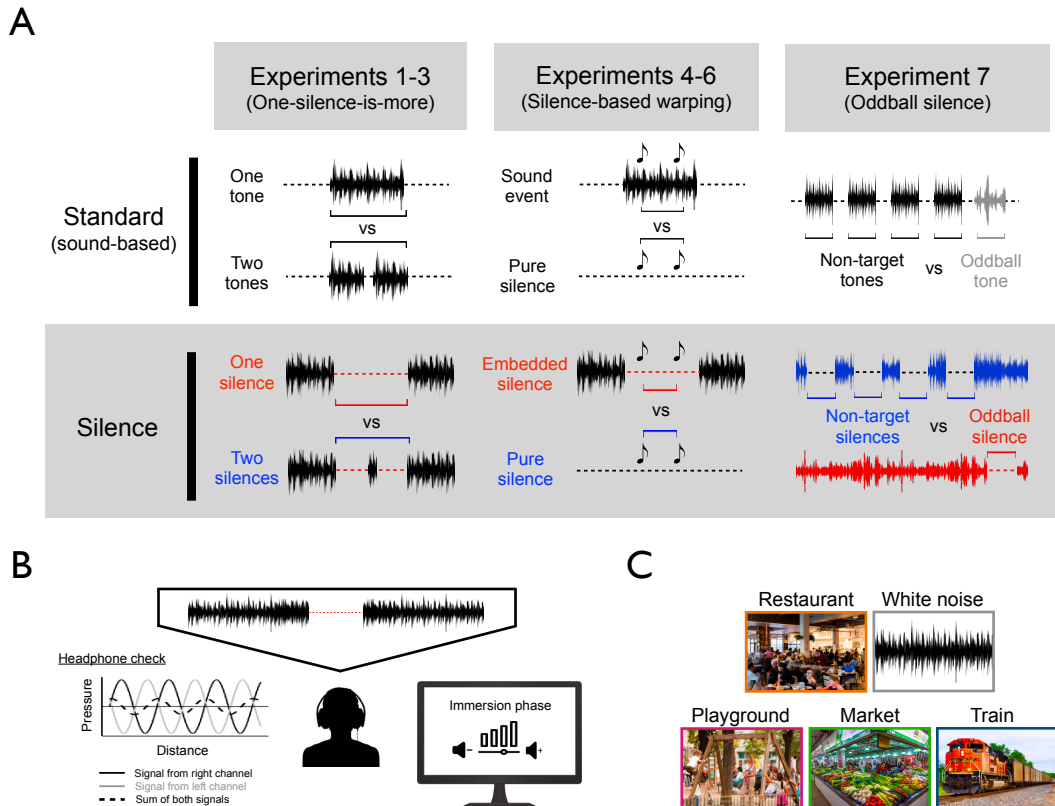


Figure 1. Experimental design and setup. (A) Overview of our ‘substitution’ approach, including the original sound-based versions of the illusions we explore here (top) and their silence-based counterparts (bottom), created by substituting silences for sounds. (B) Setup and immersion procedure for Experiments 1-7. Subjects wore headphones and underwent a check exploiting the fact that anti-phase tones sound different on headphones than loudspeakers (48). In all experiments, an ambient noise played until subjects were fully immersed in the soundscape. Periods of silence were introduced during individual trials by briefly cutting out the ambient noise. (C) Ambient noises used for Experiments 1 and 2.

Experiments 4–6: Silence-based warping

Do our results generalize to other instances of event-based temporal distortions? To test this, we created a second new auditory illusion in Experiment 4: *silence-based warping*. This paradigm was inspired by *object-based warping*, a visual illusion in which a pair of dots within an object (e.g., a rectangle) is perceived as farther apart than an equidistant pair of dots in empty space (30). Building on previous findings that object-based perceptual phenomena have event-based counterparts (25,31,32), we hypothesized that a pair of tones within an auditory event would be perceived as further apart *in time* than a pair of tones not within any auditory event. Our experiment had two conditions: *embedded silence* trials, in which subjects were immersed in ambient noise and heard tone pairs during periods of silence that interrupted the noise, and *pure silence* trials, in which subjects simply heard tone pairs in complete silence, without any ambient noise. After hearing each tone pair, subjects judged the duration between

tones with respect to a previously memorized reference duration. We predicted that embedded silences would elicit auditory event representations that would in turn dilate the perceived duration between tones (analogous to the distance between dots), while pure silences would not. Our results support this prediction — observers judged tones in embedded silences to be further apart than tones in pure silence ($t(98)=3.94$, $p<.001$; Fig. 2D).

Experiment 5 controlled for the possibility that the temporal dilation we observed in embedded silence trials was simply due to subjects being surprised or distracted by the sudden offset of ambient noise. To rule this out, we replaced pure silence trials with *surprise control* trials, which featured a brief burst of white noise at the exact time subjects would have experienced the offset of ambient noise in embedded silence trials. Subjects still judged tones in embedded silence trials to be further apart than tones in surprise control trials ($t(94)=2.76$, $p=.007$), demonstrating that silence-based warping goes beyond the influence of surprise, and further showing that an extended period of immersion (but not a brief burst of noise) is necessary to elicit event representations of silence (Fig. 2E).

Experiment 6 tested another potential confound: in embedded silence trials, subjects were reimmersed in ambient noise after hearing the tone pair, and so made their judgments in ambient noise; by contrast, subjects in pure silence trials made their judgments in silence. Could this explain our results? To address this concern, we replaced the pure silence trials from Experiment 4 with *response control* trials, in which subjects heard tone pairs in complete silence, after which they experienced the onset of ambient noise and made their judgments while immersed in noise. Subjects judged tones in embedded silence trials to be further apart than tones in response control trials ($t(88)=3.32$, $p<0.002$), showing that our results cannot be explained by differences in response conditions (Fig. 2F). Collectively, Experiments 4–6 further exemplify our ‘substitution’ strategy, revealing another event-based auditory illusion caused by periods of silence, and reinforcing our claim that perception treats such silences as genuine auditory objects.

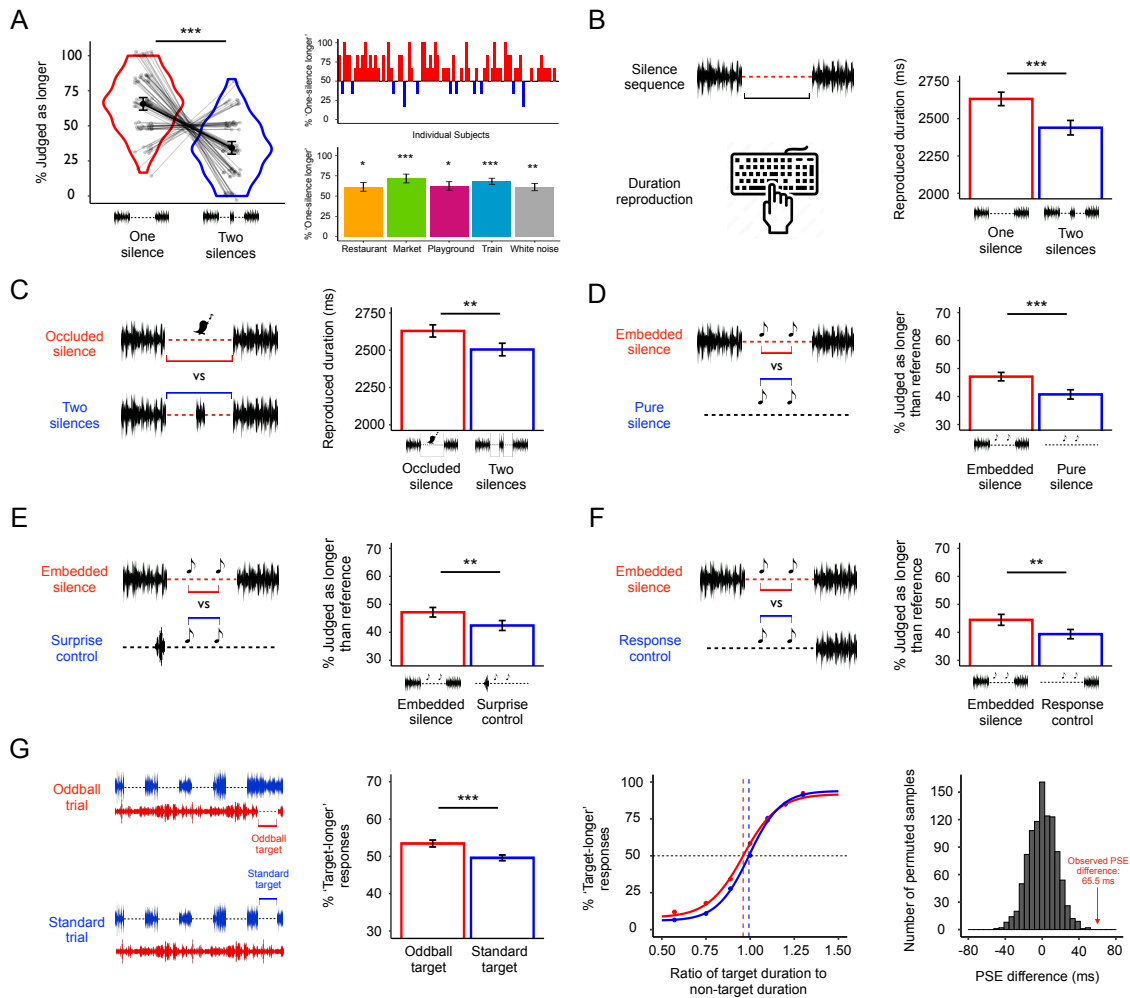


Figure 2. Results for Experiments 1-7. (A) Experiment 1 results. (Left) Percentage of 'longer' responses during equal duration trials, collapsed across subjects and ambient noise conditions. (Top right) Percentage of 'one-silence longer' responses for each subject. (Bottom right) Percentage of 'one-silence longer' responses for each ambient noise condition. Error bars represent 95% CI. (B) Experiment 2. When asked to reproduce the durations of silence sequences (rather than make forced-choice responses), a similar pattern of results emerged. Error bars represent SEM. (C) Experiment 3. (Left) Addition of an 'occluded silence' condition which contained an intervening noise that was not a resumption, but a different noise altogether (a bird chirp) played over a single continuous silence. (Right) Mean reproduced duration by sequence type. Error bars represent SEM. (D) Experiment 4. (Left) Subjects experienced two test tones in either embedded silence or pure silence, and judged whether the duration between the two test tones was longer or shorter than a previously memorized reference duration. (Right) Percentage of 'longer than reference duration' responses by silence type. Error bars represent SEM. (E) Experiment 5. (Left) To control for any surprise caused by the offset of ambient noise in embedded silence, surprise control trials included a brief burst of white noise. (Right) Percentage of 'longer than reference duration' responses by silence type. Error bars represent SEM. (F) Experiment 6. (Left) To control for differences in response conditions, response control trials included onset of the ambient noise after the test tones, during response. (Right) Percentage of 'longer than reference duration' responses by silence type. Error bars represent SEM. (G) Experiment 7. (Leftmost) On each trial, subjects heard four identical non-target silences, in which one soundtrack went silent for a fixed duration; thereafter, subjects heard a target silence, which was either a standard silence (in which the same sound went silent again), or an oddball silence (in which the other sound went silent). Subjects judged whether the target silence was longer or shorter than each of the non-targets. (Center-left) Proportion of 'target-longer' responses by target type. Error bars represent SEM. (Center-right) Psychometric curves for standard targets (blue) and oddball targets (red). (Rightmost) Empirical null distribution of PSE differences between oddball and standard conditions target from 1000 permuted samples. Red arrow indicates observed PSE difference.

Experiment 7: Oddball silences

Thus far, our experiments have investigated silences created by the complete removal of a central, salient sound. However, daily life also seemingly presents us with experiences of *partial silence*, in which one sound within a broader soundscape goes silent — such as when the bass suddenly drops out during a piece of music. Do these partial silences also elicit auditory event representations? To answer this question, we introduce a third new auditory illusion — the *oddball silence* illusion. This illusion is inspired by the auditory oddball illusion, in which a high tone that disrupts a regular sequence of low tones is perceived as longer (32,33). Experiment 7 applied our silence-substitution approach, by asking whether a novel partial silence can be perceived as ‘odd’ relative to a regular sequence of partial silences. Subjects were immersed in a soundscape comprising two distinct ‘soundtracks’ played simultaneously (e.g., a high sustained organ note and a low rumbling engine sound). On each trial, subjects heard four identical non-target silences, in which one of the two sounds went silent briefly for a fixed duration (while the other sound kept playing); thereafter, subjects heard a target silence that lasted a variable duration and had to judge whether the target silence was longer or shorter than the non-targets. Critically, the target silence was either a *standard silence*, in which the same sound that had already gone silent four times went silent again for a fifth time (e.g., four engine silences followed by a fifth engine silence), or an *oddball silence*, in which the other sound went silent (e.g., four engine silences followed by an organ silence). Strikingly, we found an analogous oddball illusion with silences — subjects judged the target silence to be longer when the target was an oddball silence than when it was a standard silence ($t(367)=3.54$, $p<.001$; Fig. 2G). This suggests not only that partial silences elicit perceptual event representations, but also that different partial silences can elicit event representations with different content. In other words, different silences ‘sound’ different.

General Discussion

The seven experiments reported above show that silences can substitute for sounds in three prominent auditory illusions caused by event representation. Across these various paradigms and phenomena, our results suggest that subjects do not simply register the durations of silent intervals but rather construct object-like representations of silences which persist through occlusion and cause temporal warping effects analogous to spatial warping in visual objects. That the same event-based illusions are elicited by both sounds and silences demonstrates that the auditory system constructs event representations of silence, just as it does with sounds.

Our findings go beyond previous work showing that the brain is sensitive to brief (≤ 50 ms) auditory gaps (34,35,22). The phenomena under study here correspond to much longer (1-5s) breaks in the auditory stream — a timescale characteristic of ordinary experiences of silence such as a dramatic pause during a speech, or the hush after an orchestral performance. Nevertheless, our results complement existing neuroscientific work demonstrating the causal role of neuronal responses in detecting auditory gaps (e.g., in mouse auditory cortex; 22); analogous neural mechanisms in humans may well play an important role in the formation of representations of silence.

Really hearing silence

The phenomena we explore here are not only evinced by statistical analyses of collected data but can also be experienced subjectively, as in our ‘demos’ of each of the illusions we report (<https://perceptionresearch.org/silence>). Moreover, the temporal distortions we observe contrast with familiar effects arising from post-perceptual mnemonic representations. For instance, the one-is-more illusion, in which one perceptual event is perceived as *longer* than two perceptual events, contrasts with the effects of event segmentation in memory, where remembered sequences comprising fewer events are judged to be *shorter* than remembered sequences (of equal objective duration) comprising more events (36). This contrast suggests that the direction of temporal distortion may help distinguish distortions of perceived duration caused by perceptual event representations from distortions of temporal judgments caused by post-perceptual event representations (i.e., reflecting later processes of compression, reconstruction, and/or retrieval). For another paradigm which finds a similar contrast, and might thus also be considered a perceptual as opposed to a memory effect, see ref 37.

Based on these considerations, we argue that silences can be genuine objects of perception. Contrary to tradition (1,6,9), we hear not only sounds, but silences. The mechanism underlying silence perception is auditory event segmentation — a process that allows perception to go beyond sensory input to track distal events. Whereas existing empirical work on auditory segmentation has focused on how periods containing acoustic information are segmented into discrete sounds or streams of sound (3,4,14-23) — reflecting the assumption that sounds are the basic units of auditory perception — our results suggest that the scope of auditory segmentation is broader: *Empty periods of time* can also be segmented to produce representations of silence.

What is the nature of these silence representations? Since representations of silence do not correspond to periods of sound, a natural way to conceive of them is as contentless event representations (a kind of empty ‘event file’; 38), or representations containing only non-acoustic temporal information. From this perspective, our findings might be seen as evidence that auditory event representations can arise even in the absence of positive acoustic content. Our results could also be understood in terms of theories of *perceptual indexing* (e.g., the FINST theory; 39), under which representations of silence might be characterized as indexes assigned to auditory absences. On these and related framings, this work points towards a broader conception of the representational processes underlying perceptual segmentation and tracking, allowing for representation and tracking of events with no positive sensory content.

Looking ahead

Our approach speaks first and foremost to the perception of particular, contrastive silences — that is, silences corresponding to the temporary absence of specific environmental sounds, such as a conversation, a musical performance, or the noise of a restaurant. These are the silences we meet with in ordinary life, and they are also among the silences deemed impossible to

perceive by the philosophical tradition which motivated our work. Importantly, these silences are nevertheless *total*: Aside from the partial silences in Experiment 7, the soundtracks in our experiments did not merely attenuate but rather went totally silent. It is a further question whether humans can also perceive *absolute* silence, as might occur during a complete lack of auditory stimulation. Given the omnipresence of internal sounds (e.g., due to blood flow, respiration, or otoacoustic emissions), absolute silence may not be physiologically possible to achieve (though see discussion in 11,13). Further work may explore the limits of our perception of silence, including the durations over which it is possible to perceive silence and precisely what kinds of sounds we can perceive as absent.

Another contribution of the present work is to introduce a general research strategy that may be used to study other forms of absence perception, including in other sensory modalities. Absence perception is challenging to study using the methods of psychophysics in part because experiences of absence typically do not reflect properties traditionally studied in perception science, such as pitch and loudness in audition, and shape and color in vision. Though there have been creative studies of conceptually related phenomena such as holes (40,41), shadows (42), and negative parts (43), our methodology overcomes this difficulty by focusing on how absences affect the perception of *temporal duration*, using established perceptual illusions known to occur with sounds and visual objects.² Recently, we have applied this methodology to studying visual absences, discovering a *one-disappearance-is-more* illusion in which a single continuous disappearance of a visual object is perceived as longer than two discrete disappearances with equal object duration (44). We thus hope that this methodology will allow perception science to further investigate the nature of absence perception across the senses.

Finally, we note that the present work adds to a burgeoning trend of collaborative engagement between philosophers and scientists, leading to real scientific progress on philosophical questions that had previously seemed empirically intractable (45). These collaborations not only explore how stubborn philosophical questions may be amenable to empirical investigation, but also highlight ways in which scientific research and theorizing may benefit from philosophical insight and inspiration (46).

² Holes, in particular, provide a potentially useful analogy with the kinds of silence we investigate here. For example, just as our silences require an auditory 'host' – the sounds which precede and succeed them – holes require a material 'host'. However, an important disanalogy is that the silences we study here are temporally (rather than spatially) extended, and so while holes and their surrounding hosts co-occur at all points in time, the sounds which surround our contrastive silences are absent during the silent period itself. Hence, while the representation of holes can be supported by synchronous visual input, representations of silence are more remarkable since they correspond to moments in time lacking in any concurrent auditory input. Though it is natural to ask whether experiences of silence correspond to related phenomena in other sense modalities, we suspect that ultimately there is no perfect analogy — and thus that it is best to regard silence on its own terms.

Materials and Methods

In this section, we provide detailed descriptions of the methods, analyses and results for each of the seven experiments reported above. All sample sizes, exclusion criteria, analyses and key experimental parameters reported here have been pre-registered. Data, analyses, stimuli and pre-registrations are publicly available at <https://osf.io/ytzxv/>. Readers can also experience all seven experiments for themselves at <https://perceptionresearch.org/silence>.

General Methods (All Experiments)

Subjects

All subjects were adults recruited from the online platform Prolific (for validation of the reliability of this subject pool, see ref 47). Each subject participated in only one experiment. Experiments 1-6 each recruited 100 subjects (600 subjects total). Experiment 7 recruited 400 subjects. All subjects provided informed consent and were compensated financially for their participation. The experiments were approved by the Homewood Institutional Review Board of Johns Hopkins University.

Stimulus Delivery

To promote immersion in the auditory stimuli, subjects in all seven experiments were required to wear headphones or earbuds for the entire experiment. To ensure this, all subjects had to pass a headphone screening procedure which consisted of six trials: On each trial, subjects heard three tones and judged which tone was quietest. One of the three tones was presented 180° out of phase across stereo channels. This tone sounds different on headphones (because each ear receives audio from only one channel) compared to loudspeakers (where both ears receive audio from both channels), thus making the task easy with headphones but prohibitively difficult with loudspeakers (for more details about this screening procedure, see ref 48). Subjects had to answer at least five out of six trials correctly before they were allowed to participate in the actual experiments. Beyond requiring headphones, we also optimized our experiment for Google Chrome (and tested it on multiple machines), and then required all subjects to use Google Chrome as well. If any subject initiated the experiment using a non-Chrome browser, they were required to switch to Chrome before being allowed to participate.

Experiment 1: One-silence-is-more (comparison)

This experiment introduces the *one-silence-is-more illusion*, which was adapted from the (sound-based) one-is-more illusion, in which a single continuous tone is perceived as longer than a sequence comprising two tones with the same total duration (25). We inverted the one-is-more illusion by substituting silences for the tones: rather than ask whether one long sound is perceived as longer than a sequence comprising two short sounds, we asked whether one long silence is perceived as longer than a sequence comprising two short silences. The pre-registration for this experiment is available at https://aspredicted.org/KY5_8NS.

Stimuli and procedure

For the entire duration of the experiment (except for the silence sequences occurring on each trial), subjects were immersed in ambient noise. The ambient noise track was randomly chosen for each subject from five possible options: Restaurant, Train, Playground, Market, and White Noise. The ambient noise soundtracks were obtained from online sound repositories and were looped using Audacity audio editor (version 2.4.2; <https://www.audacityteam.org/>) to fit the duration of the experiment. At the start of the experiment, subjects were told to adjust the volume of the ambient noise until they felt sufficiently immersed in the soundscape (for instance,

subjects who heard the restaurant noise were told to “adjust your volume until it sounds like you’re actually sitting in a busy restaurant”).

On each trial of the experiment, two sequences of silence were presented successively. In one of the sequences, the ambient noise was interrupted by a single continuous silence (the one-silence sequence). In the other sequence, the ambient noise was interrupted by two discrete silences separated by a brief period of noise resumption (the two-silences sequence). The noise resumption in two-silences sequences was always $1/9^{\text{th}}$ of the total duration of the sequence. The order of sequence presentation was counterbalanced within subjects. Each trial began with an immersion period of 5s, after which subjects heard an announcer say “one”, followed 1s later by the presentation of the first silence sequence. 3s after the end of the first sequence, subjects heard the announcer say “two”, followed 1s later by the presentation of the second silence sequence. 2s after the end of the second sequence, a prompt appeared on the screen asking the subject to press the “1” key if the first sequence was longer or press the “2” key if the second sequence was longer. The next trial began once the subject responded. During the instruction phase, subjects were explicitly told (using both text and graphical illustrations) that their task was to compare the durations of the whole sequences, not just the durations of the silences. To ensure that subjects understood their task, they were required to complete two ‘easy’ practice trials (in which one sequence was 1.125s long and the other sequence was 4.5s long) during the instruction phase. If they answered wrongly, they were reminded of the instructions and had to attempt the same practice trial again. Subjects had to answer both practice trials correctly before they could start the actual experiment. If they failed any practice trial three times, subjects were disqualified from the study.

In the actual experiment, there were three trial types, appearing in the experiment with equal frequency and in random order: trials in which the one-silence sequence was longer than the two-silences sequence (single-longer trials), trials in which the two-silences sequence was longer than the one-silence sequence (double-longer trials), and trials in which both sequences were equally long (equal-duration trials). Each trial also belonged to one of three duration categories: short (1-1.25s), medium (2-2.5s) and long (4-5s). Within each category, there were three possible sequence durations: 1s, 1.125s, 1.25s; 2s, 2.25s, 2.5s; 4s, 4.5s, 5s. The exact sequence durations for each trial were determined using a pseudo-random algorithm that was constrained by the trial’s type and duration category.

Each subject completed 18 experimental trials (3 trial types x 3 duration categories x 2 sequence presentation orders) and 2 catch trials in which one sequence was sampled from the short duration category and the other sequence was sampled from the long duration category. The 20 total trials were presented in random order.

Analyses and results

In accordance with our pre-registered analysis plan, we excluded subjects if they responded incorrectly to at least one catch trial. 85 subjects remained after exclusion.

On equal-duration trials, subjects consistently chose the one-silence sequence as longer than the two-silences sequence ($t(84)=6.96$, $p<.001$). Strikingly, the proportion of ‘one-silence longer’ responses in our experiment was almost identical to the proportion of ‘one-tone longer’ responses in the original one-is-more illusion (0.66 in both cases), suggesting that the same event-based phenomenon is at play in both experiments. Moreover, this bias toward ‘one-silence longer’ responses on equal-duration trials remained statistically significant within each duration category and for each soundtrack, demonstrating the generality of the one-silence-is-more illusion across multiple timescales and ambient noises.

To ensure that subjects did not simply judge the one-silence sequence to be longer because it contained more silence than the two-silences sequence, we also analyzed double-longer trials in which both sequences had equal durations of silence, and found that subjects still judged the one-silence sequence as longer than than the two-silences sequence ($t(78)=2.46$, $p=.016$).

Experiment 2: One-silence-is-more (reproduction)

In this experiment, we replicated the one-silence-is-more illusion with a different behavioral measure. Subjects in Experiment 2 experienced only one sequence on each trial, and were asked to reproduce the duration of the single sequence they heard. The pre-registration for this experiment is available at https://aspredicted.org/KFY_KZH.

Stimuli and procedure

The stimuli and procedure for Experiment 2 were identical to Experiment 1 except as specified below.

Instead of experiencing two sequences and comparing their durations (as in Experiment 1), subjects in Experiment 2 experienced only one sequence on each trial – either a one-silence sequence or a two-silences sequence. Each trial began with an immersion period of 6s, followed by the presentation of a silence sequence. 2s after the end of the silence sequence, a prompt appeared on the screen asking the subject to reproduce the duration of the sequence they had just heard by holding down the spacebar. If subjects pressed the spacebar before the prompt appeared, they received a warning message reminding them to press the spacebar only after the prompt appears, and had to redo the trial. If subjects pressed the spacebar more than once after the prompt appeared, they received a warning message reminding them to hold down the spacebar only once to reproduce the duration of the sequence they just heard, after which they advanced to the next trial.

During the instruction phase, subjects were explicitly told that their task was to reproduce the duration of the whole sequence, not just the duration of the silences. To ensure that subjects understood the task, they were required to complete two practice trials during the instruction phase: the first trial was a 2.5s one-silence sequence and the second trial was a 2.5s two-silences sequence. If a subject's response on a practice trial was 1s longer or shorter than the actual sequence duration, the subject failed the practice trial, was reminded of the instructions, and had to redo the trial. Subjects were allowed to start the experiment only after they passed both practice trials. If they failed any practice trial three times, subjects were disqualified from the study.

In the actual experiment, there were three possible sequence durations: short (1.25s), medium (2.5s) and long (5s). Each subject completed 12 experimental trials (2 conditions x 3 durations x 2) presented in random order.

Analyses and results

In accordance with our pre-registered analysis plan, subjects who had at least two trials in which the reproduced duration was 75% longer or shorter than the actual event duration were excluded. 97 subjects remained after these exclusions. After excluding these subjects, we also excluded trials in which the reproduced duration was less than 250ms (this was to exclude trials in which subjects accidentally pressed the spacebar), as well as trials in which subjects pressed the spacebar more than three times before the prompt appeared, and trials in which subjects pressed the spacebar more than once after the prompt appeared. This led to the exclusion of 8 additional trials in total across all subjects.

Mean reproduced duration was significantly longer for one-silence sequences than two-silences sequences ($t(96)=4.73$, $p<.001$), demonstrating that the one-silence-is-more illusion replicates with a duration reproduction task.³

Instructions control

To further ensure that the effects reported in this experiment do not reflect misunderstanding of the task (e.g., reproducing the total time spent in silence, or just one of the silences in the two-silences condition, instead of the total durations of the entire sequences), we also ran a follow-up experiment that was identical to Experiment 2 except that, at the conclusion of the experimental session, all subjects answered a debriefing question asking them to verify the task instructions. The debriefing question presented subjects with a sample trial (a two-silences trial with two 1s long silences and a 0.2s resumption) and subjects were asked to indicate which of three options was the correct reproduction length. The three options were: (a) 1 second (the length of one silence), (b) 2 seconds (the total length of both silences excluding the resumption) and (c) 2.2 seconds (the total length of both silences including the resumption), appearing in a random order. The pre-registration for this experiment is available at https://aspredicted.org/DHS_8CZ.

Even excluding any subject who failed to answer this question correctly (i.e., considering only those subjects who correctly verified the instructions), the one-silence-is-more illusion emerged, $t(80)=4.05$, $p<.001$. We thank a reviewer for comments that led to this follow-up experiment.

Experiment 3: One-silence-is-more (contrast control)

This experiment tested a potential alternative explanation of our results, namely that the one-silence-is-more illusion is not caused by event representations, but instead arises due to the attentional effects of the intervening noise in the two-silences sequence (e.g., if subjects were distracted by the intervening noise). To address this confound, we added a third sequence type: the '*occluded-silence*' sequence, which contained an intervening noise that was not a resumption of ambient noise, but a different noise altogether (a bird chirping) which sounds like it's being played over a single continuous silence. The pre-registration for this experiment is available at https://aspredicted.org/CHH_PK8.

Stimuli and procedure

The stimuli and procedure for Experiment 3 were identical to Experiment 2 except with the addition of occluded-silence trials and the use of only white noise as ambient noise.

Occluded-silence trials were identical to two-silences trials except that the intervening noise was not a resumption of ambient noise, but a bird chirping noise. Three bird chirping clips were made, each matching the duration of ambient-noise resumption in the short, medium and long silence sequences. The bird chirping noises and corresponding ambient-noise resumptions were equated for loudness using the Loudness Normalization function in Audacity (version 2.4.2; <https://www.audacityteam.org/>), which normalizes the loudness of the two tracks to a fixed level of -16.0 LUFS (the homogeneity of white noise allowed for exact equation).

³ Mean reproduced duration was also significantly longer for one-silence sequences than two-silences sequences when the analysis was repeated without excluding any subjects or trials ($t(99)=4.68$, $p<.001$). This secondary analysis was done to mirror the lack of exclusion criteria in the original experiment that demonstrated the one-is-more illusion with sounds (25).

During the instruction phase, subjects had to complete an occluded-silence practice trial in addition to the two practice trials specified in Experiment 2. Each subject completed 18 experimental trials (3 conditions x 3 durations x 2) presented in random order.

Analyses and results

14 subjects and 8 additional trials were excluded via our pre-registered exclusion criteria, which were identical to the exclusion criteria of Experiment 2.

As in Experiment 2, mean reproduced duration was longer for one-silence trials than for two-silences trials ($t(81)=4.12, p<.001$), once again replicating the one-silence-is-more illusion. Crucially, and consistent with our hypothesis, mean reproduced duration was also longer for occluded-silence trials than two-silences trials, ($t(85)=2.92, p<.005$). This shows that the one-silence-is-more illusion occurs even when both sequences contain equally loud intervening noises.⁴

Experiment 4: Silence-based warping

In this experiment, we introduce a second new silence illusion — *silence-based warping*. Our silence-based warping paradigm was inspired by object-based warping, a visual illusion in which a pair of dots within an object (e.g., a rectangle) is perceived as farther apart than an equidistant pair of dots in empty space (30). Here, we demonstrate an event-based analog of object-based warping, in which a pair of *tones* within a *silence event* is perceived as further apart *in time* than a pair of tones not within any auditory event. The pre-registration for this experiment is available at https://aspredicted.org/VB9_VV1.

Stimuli and procedure

During the instruction phase, subjects listened to a pair of reference tones played successively and were asked to memorize the duration between the tones; we refer to this duration as the *reference duration* (all tones used in this experiment were 493.88Hz and lasted 200ms). Subjects had to play the reference tones at least three times during the instruction phase to memorize it. The reference duration was chosen randomly for each subject from three options (1s, 1.5s, 2s) and was constant within each experiment.

The experiment consisted of 8 blocks. Before the start of each block, subjects had to play the reference tones three more times to remind themselves of the reference duration. There were two block types, corresponding to the two experimental conditions: *embedded silence* blocks and *pure silence* blocks.

In embedded silence blocks, subjects were immersed in the ambient noise of a restaurant (the same restaurant track as in Experiment 1). The ambient noise played for the whole block, except during the silences. Each trial began with 4s of immersion time, before the ambient noise cut off. After a silent buffer interval (see below for the possible durations of this interval), two test tones were played successively. We refer to the duration between test tones as the *test duration*. After another identical buffer interval, the ambient noise resumed, followed 1.5s later by a prompt appearing on-screen asking the subject to press the “L” key if the test duration was longer than the reference duration, or press the “S” key if the test duration was shorter than the reference duration. The next trial began once the subject responded.

⁴ Interestingly, mean reproduced duration for occluded-silence trials was marginally shorter than for one-silence trials than occluded-silence trials ($t(85)=1.99, p=.0498$), suggesting (perhaps unsurprisingly) that attentional effects of the intervening noise may have a slight effect on reproduced duration. However, the primary comparison between occluded-silence trials and two-silences trials demonstrates that the attentional effects of intervening noise cannot fully explain the one-silence-is-more illusion.

Pure silence blocks were identical to embedded silence blocks except that subjects were not immersed in any ambient noise. Each trial began with 4s of complete silence, followed by a silent buffer interval. Thereafter, a pair of test tones was played successively, followed by another identical buffer interval. 1.5s later, a prompt appeared on-screen asking the subject to press the “L” key if the test duration was longer than the reference duration, or press the “S” key if the test duration was shorter than the reference duration. The next trial began once the subject responded.

To ensure that subjects understood their task, they were required to complete four ‘easy’ practice trials during the instruction phase in the following order: 1) pure silence trial with 4s test duration, 2) pure silence trial with 0.5s test duration, 3) embedded silence trial with 4s test duration, 4) embedded silence trial with 0.5s test duration. If a subject answered any practice trial wrongly, they were reminded of the instructions and had to redo the trial. Subjects were allowed to begin the experiment only after they passed all four practice trials. If they failed any practice trial three times, subjects were disqualified from the study.

Each experimental block had seven experimental trials, each with a different test duration. The test durations were ratios (0.85, 0.90, 0.95, 1.0, 1.05, 1.10, 1.15) of the reference duration. Trials in each block were presented in random order. There were two possible buffer intervals (0.5s and 1s). The buffer intervals were constant within each block and counterbalanced between blocks. In addition to the experimental trials, there were also two easy catch trials, one with a 4s test duration and the other with a 0.5s test duration. One catch trial was randomly inserted into an embedded silence block, while the other catch trial was randomly inserted into a pure silence block. Each subject completed 56 experimental trials (2 conditions x 4 blocks x 7 test durations) and 2 catch trials.

Analyses and results

In accordance with our pre-registered exclusion criteria, we excluded any subject who failed at least one catch trial. 99 subjects remained after exclusions. No additional trials were excluded.

The following analysis collapses across all reference durations, test duration ratios and buffer intervals. Subjects were more likely to judge the test duration to be longer than the reference duration in embedded silence trials than in pure silence trials ($t(98)=3.94$, $p<.001$), suggesting that embedded silences elicit event representations which warp perceived duration while pure silences do not.

Experiment 5: Silence-based warping (surprise control)

In this experiment, we controlled for the possibility that the temporal dilation we observed in embedded silence trials was simply due to subjects being surprised or distracted by the sudden offset of ambient noise. To rule this out, we replaced pure silence trials with *surprise control* trials, which featured a brief burst of white noise at the exact time subjects would have experienced the offset of ambient noise in embedded silence trials. The pre-registration for this experiment is available at https://aspredicted.org/YXX_D92.

Stimuli and procedure

The stimuli and procedure for Experiment 5 were identical to Experiment 4, except that all pure silence trials (including catch trials) were replaced by surprise control trials.

The only difference between pure silence trials and surprise control trials was that pure silence trials began with 4s of complete silence, whereas surprise control trials began with 3.75s of

complete silence followed by a 0.25s burst of white noise. The white noise was timed to offset at exactly the same time as the offset of the restaurant noise in embedded silence trials. The white noise and restaurant tracks were equated for perceived loudness using the Loudness Normalization function in Audacity (version 2.4.2; <https://www.audacityteam.org/>), which normalizes the loudness of the two tracks to a fixed level of -16.0 LUFS.

Analyses and results

4 subjects were excluded because they failed at least one catch trial, and 1 additional subject was excluded because they did not contribute a full dataset. 95 subjects remained after exclusions.

Subjects were more likely to judge the test duration to be longer than the reference duration in embedded silence trials than in surprise control trials ($t(94)=2.76$, $p=.007$), showing that silence-based warping is not solely due to the effects of surprise or distraction.

Experiment 6: Silence-based warping (response control)

In this experiment, we addressed another potential confound in our silence-based warping experiments. In Experiment 4, subjects in embedded silence trials were reimmersed in ambient noise after hearing the tone pair, and so made their judgments in ambient noise; by contrast, subjects in pure silence trials made their judgments in silence. To test whether our earlier results were caused by this difference in response conditions, we replaced pure silence trials with *response control* trials, in which subjects heard tone pairs in complete silence, after which they experienced the onset of ambient noise and made their judgments while immersed in noise. The pre-registration for this experiment is available at https://aspredicted.org/PBH_KP3.

Stimuli and procedure

The stimuli and procedure for Experiment 6 were identical to Experiment 4, except that all pure silence trials (including catch trials) were replaced by response control trials.

Response control trials were identical to pure silence trials, except that after the second buffer interval (which came after the test tones), the restaurant noise started playing, and subjects were immersed in restaurant noise until they responded.

Analyses and results

8 subjects were excluded because they failed at least one catch trial, and 3 additional subjects were excluded because they did not contribute a full dataset. 89 subjects remained after exclusions.

Subjects were more likely to judge the test duration to be longer than the reference duration in embedded silence trials than in response control trials ($t(88)=3.32$, $p<0.002$), showing that our results in Experiment 4 are not due to the difference in response conditions between embedded silence trials and pure silence trials.

Experiment 7: Oddball silence

In this experiment, we asked whether partial silences (in which one sound in a broader soundscape goes silent) elicit auditory event representations. To answer this question, we introduce a third silence illusion: the *oddball silence illusion*. Our oddball silence paradigm was inspired by the auditory oddball illusion, in which a high tone that disrupts a regular sequence of low tones is perceived as longer (32,33). Our subjects were immersed in a soundscape comprising two distinct 'soundtracks' played simultaneously. Instead of hearing a regular sequence of sounds, they heard a regular sequence comprising four identical partial silences in

which one of the two soundtracks went silent. Thereafter, subjects heard a target silence, which could either be *standard silence*, in which the same soundtrack that had gone silent four times went silent again for a fifth time, or an *oddball silence*, in which the other soundtrack went silent. We predicted that oddball targets would be perceived as longer than standard targets. The pre-registration for this experiment is available at https://aspredicted.org/ZFS_WX1.

Stimuli and procedure

Throughout the experiment (except for the partial silences during each trial), subjects were immersed in a soundscape comprising two contrasting soundtracks. The soundtracks were constant for each subject, but varied between subjects. One soundtrack was either a rushing waterfall or a rumbling engine, and the other soundtrack was either an organ playing a sustained note, or a violin playing vibrato. This gave rise to 4 soundtrack pairs (waterfall/organ, waterfall/violin, engine/organ, engine/violin), with 100 subjects experiencing each soundtrack pair. These soundtracks were obtained from online sound repositories and were looped using Audacity (version 2.4.2; <https://www.audacityteam.org/>) to fit the duration of the experiment. At the start of the experiment, subjects were told to adjust the volume of the soundscape until they felt sufficiently immersed.

Each trial in the experiment began with 4s of immersion in the soundscape, followed by four successive non-target silences in which one of the two sounds went silent briefly, before resuming. Non-target silences were always 2s long, and the intervals between silences were also 2s. After the fourth non-target silence and an inter-silence interval of 2s, a target silence was presented. The target silence could either be a *standard silence*, in which the sound that had gone silent four times went silent again for the fifth time (e.g., four engine silences followed by a fifth engine silence), or an *oddball silence*, in which the other sound went silent (e.g., four engine silences followed by an organ silence). There were 7 possible target silence durations: 1.4s, 1.6s, 1.8s, 2.0s, 2.2s, 2.4s, 2.6s.

The sound that went silent in the non-target silences was chosen pseudorandomly for each trial. The pseudo-random selection algorithm ensured that, for a given subject and target duration, the non-target silences were identical across both standard and oddball conditions, thus allowing for within-subjects comparisons.

In addition to the auditory stimuli, there was a counter on screen to help subjects keep track of the number of silences. During the immersion time at the start of the trial, the counter displayed “_”. At the onset of each silence, the counter changed to reflect the number corresponding to the current silence (e.g., at the onset of the first silence, the counter changed to show “1”; at the onset of the target silence, the counter changed to show “5”). 1s after the end of the target silence, a prompt appeared on the screen asking subjects to press the “L” key if the target silence was longer than each of the non-targets, or press the “S” key if the target silence was shorter than each of the non-targets. The next trial began once the subject responded.

To ensure that subjects understood the task, they were required to complete two ‘easy’ practice trials during the instruction phase. The first practice trial was a standard silence trial with 4s target duration and the second practice trial was an oddball silence trial with 0.5s target duration. If subjects failed a practice trial, they were reminded of the instructions and had to redo the trial. Subjects were only allowed to start the experiment after they passed both practice trials. If they failed any practice trial three times, subjects were disqualified from the study.

Each subject completed 14 experimental trials (2 conditions x 7 target durations), and 2 catch trials, one with a target duration of 0.5s, and the other with a target duration of 4s. These 16 trials were presented in random order.

Analyses and results

In accordance with our pre-registered exclusion criteria, we excluded subjects that failed at least 1 catch trial. 32 subjects were excluded, leaving 368 subjects.

Collapsing across all target durations and soundtrack pairs, we found that subjects were more likely to judge the target silence as longer than non-targets when the target was an oddball silence than when the target was a standard silence ($t(367)=3.54$, $p<.001$). The same trend held for each soundtrack pair (although our experiment was not adequately powered for the effects within each soundtrack pair subset to attain statistical significance in every case).⁵

To quantify the degree of temporal distortion in the oddball silence illusion, we computed the Point of Subjective Equality (PSE) separately for each condition. To do this, we first calculated the mean proportion of 'target longer' responses at each target duration, collapsing across subjects. Next, we fitted a cumulative normal function that predicts the probability of the 'target longer' response from target duration (the model we fitted was adapted from ref 49). The PSE is the target duration at which the predicted probability of responding 'target longer' is 0.5. The PSE for standard trials was 1986.4ms, while the PSE for oddball trials was 1920.8ms. We used permutation testing to determine statistical significance. More specifically, we tested the observed PSE difference against an empirical null distribution constituted by 1000 permuted samples. Each permuted sample was created by shuffling the mapping between condition labels and 'target longer/shorter' responses *within subjects* (so all permuted samples had the same nested structure as our actual sample). The observed PSE difference (65.5ms) lies outside of the 2.5 and 97.5 quantiles of the empirical null distribution ([-37.8ms, 37.5ms]), demonstrating that the PSE difference between conditions is statistically significant. We also calculated a 'p-value', which here is just the proportion of samples in the empirical null distribution with PSE differences that are greater than our observed PSE difference. This proportion is 0 (our observed PSE difference is 65.6ms, while the highest permuted PSE difference is 65.0ms).

Acknowledgments

For helpful feedback and discussion, we gratefully acknowledge audiences at the Vision Sciences Society, Society for Philosophy and Psychology, Southern Society for Philosophy and Psychology, Princeton University, Neural Mechanisms, Washington University in St. Louis, University of Southern California, City University of New York, University of California-Los Angeles, and members of the Johns Hopkins University (JHU) Perception & Mind Lab. For consultation on statistical analyses, we thank Jeff Bowen. We also thank the editor and two reviewers at PNAS for helpful comments that led to significant improvements to the paper.

⁵ A secondary analysis breaking subjects and trials into subgroups suggested that the silence oddball effect was primarily driven by 'background' silences (e.g., the engine going silent after four organ silences) as opposed to 'foreground' silences (e.g., the organ going silent after four engine silences). Similar patterns arise in the oddball literature, where certain stimuli (e.g., an expanding circle) give rise to bigger oddball effects than other stimuli (e.g., a stationary circle; see ref 33).

References

1. A. A. Luce, T. E. Jessop, Eds., *The Works of George Berkeley, Bishop of Cloyne*, (Thomas Nelson, 1948-1957).
2. C. J. Shields, Ed., *Aristotle: De Anima* (Oxford University Press, 2016).
3. A. S. Bregman, *Auditory Scene Analysis: The Perceptual Organization of Sound* (MIT Press, 1994).
4. M. Kubovy, Auditory and visual objects. *Cognition* **80**, 97-126 (2001).
5. M. Nudds, "Audition" in *Oxford Handbook of Philosophy of Perception*, M. Matthen, Ed. (Oxford University Press, 2015), pp. 274-293.
6. G. J. Warnock, *Berkeley* (University of Notre Dame, 1983).
7. R. M. Warren, *Auditory Perception: A New Synthesis*. (Elsevier, 2013).
8. J. Schnupp, I. Nelken, A. J. King, *Auditory Neuroscience: Making Sense of Sound* (MIT Press, 2011).
9. B. O'Shaughnessy, *Consciousness and the World* (Oxford University Press, 2000).
10. D. Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (MIT Press, 2010).
11. R. Sorenson, *Seeing Dark Things: The Philosophy of Shadows* (Oxford University Press, 2008).
12. B. Skrzypulec, Silence Perception and Spatial Content. *Australas. J. Philos.* **100**, 1-15 (2021).
13. I. Phillips, "Hearing and hallucinating silence" in *Hallucination: Philosophy and Psychology*, F. Macpherson, D. Platchias, Eds. (MIT Press, 2013).
14. L. Lisker, A. S. Abramson, A cross-language study of voicing in initial stops: Acoustical measurements. *Word* **20**, 384-422 (1964).
15. M. F. Dorman, L. J. Raphael, A. M. Liberman, Some experiments on the sound of silence in phonetic perception. *J. Acoust. Soc. Am.* **65**, 1518-1532 (1979).
16. S. S. Reich, Significance of pauses for speech perception. *J. Psycholinguist. Res.* **9**, 379-389 (1980).
17. D. Duez, Perception of silent pauses in continuous speech. *Lang. Speech* **28**, 377-389 (1985).
18. J. W. Peter, The beginnings of word segmentation in English-learning infants. *Cogn. Psychol.* **39**, 159-207 (1999).
19. M. Boltz, Rhythm and "good endings": Effects of temporal structure on tonality judgments. *Percept. Psychophys.* **46**, 9-17 (1989).
20. J. G. Martin, W. Strange, The perception of hesitation in spontaneous speech. *Percept. Psychophys.* **3**, 427-438 (1968).
21. C. Kopp-Scheinflug, et al., The sound of silence: ionic mechanisms encoding sound termination. *Neuron* **71**, 911-925 (2011).
22. A. P. Weible, et al., Perceptual gap detection is mediated by gap termination responses in auditory cortex. *Curr. Biol.* **24**, 1447-1455 (2014).
23. H. C. Hughes et al., Responses of human auditory association cortex to the omission of an expected acoustic event. *Neuroimage* **13**, 1073-1089 (2001).

24. A. Lotto, L. Holt, Psychology of auditory perception. *Wiley Interdiscip. Rev. Cogn. Sci.* **2**, (2011).
25. S. Yousif, B. J. Scholl, The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time. *Cognition* **185**, 121-130 (2019).
26. R. M. Warren, J. M. Wrightson, J. Puzos, Illusory continuity of tonal and infratonal periodic sounds. *J. Acoust. Soc. Am.* **84**, 1338-1342 (1988).
27. G. E. Thinès, A. E. Costall, G. E. Butterworth, Eds., (1991). *Michotte's Experimental Phenomenology of Perception* (Lawrence Erlbaum Associates, 1991).
28. B. J. Scholl, Z. W. Pylyshyn, Tracking multiple items through occlusion: Clues to visual objecthood. *Cogn. Psychol.* **38**, 259-290 (1999).
29. J. Munton, How to see invisible objects. *Noûs* **56**, 343-365 (2022).
30. T. J. Vickery, M. M. Chun, Object-based warping: An illusory distortion of space within objects. *Psychol. Sci.* **21**, 1759-1764 (2010).
31. J. De Freitas, B. M. Liverence, B. J. Scholl, Attentional rhythm: A temporal analogue of object-based attention. *J. Exp. Psychol.* **143**, 71-76 (2014).
32. T. Birngruber, H. Schröter, R. Ulrich, Duration perception of visual and auditory oddball stimuli: Does judgment task modulate the temporal oddball effect?. *Atten. Percept. Psychophys.* **76**, 814-828 (2014).
33. P. U. Tse, J. Intriligator, J. Rivest, P. Cavanagh, Attention and the subjective expansion of time. *Percept. Psychophys.* **66**, 1171-1189 (2004).
34. K. B. Snell, D. R. Frisina, Relationships among age-related differences in gap detection and word recognition. *J. Acoust. Soc. Am.* **107**, 1615-1626 (2000).
35. S. W. Threlkeld, S. C. Penley, G. D. Rosen, R. H. Fitch, Detection of silent gaps in white noise following cortical deactivation in rats. *Neuroreport* **19**, 893-898 (2008).
36. R. A. Block, M. A. Reed, Remembered duration: Evidence for a contextual-change hypothesis. *J. Exp. Psychol. Learn. Mem. Cogn.* **4**, 656-665 (1978).
37. B. M. Liverence, B. J. Scholl, Discrete events as units of perceived time. *J. Exp. Psychol. Hum. Percept. Perform.* **38**, 549-554 (2012).
38. B. Hommel, Event files: Feature binding in and across perception and action. *Trends Cogn. Sci.* **8**, 494-500 (2004).
39. Z. Pylyshyn, The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition* **32**, 65-97 (1989).
40. A. R. Albrecht, A. List, L. C. Robertson, Attentional selection and the representation of holes and objects. *J. Vis.* **8**, 8 (2008).
41. T. S. Horowitz, Y. Kuzmova, Can we track holes? *Vision Res.* **51**, 1013-1021 (2011).
42. R. A. Rensink, P. Cavanagh, The influence of cast shadows on visual search. *Perception* **33**, 1339-1358 (2004).
43. E. J. Green, M. Singh, J. Feldman, Apparent Motion of Negative Parts. *J. Vis.* **15**, 524-524 (2015).
44. R. Z. Goh, I. Phillips, C. Firestone, Seeing nothing happening: Moments of absence as perceptual events. *Journal of Vision* **22**, 3574-3574 (2022).
45. J. Morales, C. Firestone, Philosophy of perception in the psychologist's laboratory. *Curr. Dir. Psychol. Sci.* (In press).

46. L. Laplane, *et al.*, Why science needs philosophy. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 3948-3952 (2019).
47. E. Peer, L. Brandimarte, S. Samat, A. Acquisti, Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *J. Exp. Soc. Psychol.* **70**, 153-163 (2017).
48. K. J. P. Woods, M. H. Siegel, J. Traer, J. H. McDermott, Headphone screening to facilitate web-based auditory experiments. *Atten. Percept. Psychophys.* **79**, 2064-2072 (2017).
49. D. Odic, M. E. Libertus, L. Feigenson, J. Halberda, Developmental change in the acuity of approximate number and area representations. *Dev. Psychol.* **49**, 1103-1112 (2013).