Journal of Experimental Psychology: Applied
2002, Vol. 8, No. 1, 17–25

1076-898X/02/$5.00 DOI: 10.1037//1076-898X.8.1.17

Pitch and Loudness Interact in Auditory Displays: Can the Data Get Lost in the Map?

John G. Neuhoff
The College of Wooster

Gregory Kramer
Metta Foundation

Joseph Wayand
Kent State University

Many auditory displays use acoustic attributes such as frequency, intensity, and spectral content to represent different characteristics of multidimensional data. This study demonstrated a perceptual interaction between dynamic changes in pitch and loudness, as well as perceived asymmetries in directional acoustic change, that distorted the data relations represented in an auditory display. Three experiments showed that changes in loudness can influence pitch change, that changes in pitch can influence loudness change, and that increases in acoustic intensity are judged to change more than equivalent decreases. Within a sonification of stock market data, these characteristics created perceptual distortions in the data set. The results suggest that great care should be exercised when using lower level acoustic dimensions to represent multidimensional data.

The use of sound to present information and data is becoming increasingly common in a diverse array of settings that range from the classroom to the operating room. An often-cited example of a particularly successful auditory display is that of the Geiger counter, which indicates increased radioactivity with an increase in the density of its temporal acoustic pattern. The Geiger counter and other examples of auditory display have been in use since at least the early 1900s, but the recent technological revolution has made the presentation of information with sound even more economical, more effective, and more widespread (Kramer et al., 1999). Applied settings that use auditory display are diverse. Modern applications include providing target and threat location warnings to fighter pilots (Bronkhorst, Veltman, & van Breda, 1996; McKinley & Ericson, 1997; McKinley, Ericson, & D'Angelo, 1994), evaluating the structural integrity of large bridges (Valenzuela, Sansalone, Krumhansl, & Streett, 1997), and even guiding the manipulation of surgical instruments during brain surgery (Wegner, 1998).

A more specific type of auditory display technique called *sonification* involves "the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation" (Kramer et al., 1999, p. 3). Sonification is a technique for data display that typically involves mapping multivariate data sets onto acoustic parameters to sonically represent the data. In some cases, sonification may be preferable to data visualization, particularly in situations in which large numbers of changing variables or temporally complex information must be monitored (Kramer, 1994a). Alternatively, sonification can be used to augment data visualization techniques by providing informative redundancies in data representation that can enhance user performance. Applications that use sonification are also wide ranging and include monitoring data in complex work environments such as anesthesiology stations and factory production controls (Fitch & Kramer, 1994; Gaver, Smith, & O'Shea, 1991), analyzing seismology data (Hayward, 1994; Saue & Fjeld, 1997; Speeth, 1961), providing data display for the visually impaired (Flowers, Buhman, & Turnage, 1997; Flowers & Hauer, 1995; Lunney & Morrison, 1981; Turnage, Bonebright, Buhman, & Flowers, 1996), and even monitoring the oscillation of subatomic particles in quantum physics (Pereverzev, Loshak, Backhaus, Davis, & Packard, 1997).

Although the use of sonification and auditory display has increased, systematic research into the design and evaluation of these techniques has not kept pace. This is perhaps due to the difficulties in conducting the interdisciplinary research required for developing and evaluating auditory displays (Kramer et al., 1999). Nevertheless, from both an applied and theoretical perspective, a key issue in the development of effective auditory display and sonification techniques is optimizing the degree of match between the intended sonic representation of information and the perceptual experience of that information by the listener. In many sonifications, auditory variables such as loudness, pitch, and timbre are used isomorphically to represent data variables. Often two or more data variables are represented by changing two or more of these auditory parameters within one auditory stream (Kramer, 1994b). In some respects, the utility of this type of representation seems

---

John G. Neuhoff, Department of Psychology, The College of Wooster; Gregory Kramer, Metta Foundation, Portland, Oregon; Joseph Wayand, Department of Psychology, Kent State University.

This article was supported by National Science Foundation Grant 9905266 awarded to John G. Neuhoff. We thank Bruce Walker and John Flowers for helpful comments on an earlier version of this article.

Correspondence concerning this article should be addressed to John G. Neuhoff, Department of Psychology, The College of Wooster, Wooster, Ohio 44691. E-mail: jneuhoff@wooster.edu

entirely reasonable. Frequency, intensity, and spectral characteristics are easily specified, and can be mapped onto data sets relatively simply. For example, simultaneous changes in frequency and spectrum have been used to represent geological and geophysical data in gas and oil explorations (Barrass & Zehner, 2000). Changes in frequency, intensity, and oscillation rate have been used in exploratory data analysis (Hermann, Meinicke, & Ritter, 2000).

However, a potential problem with mapping data variables onto acoustic parameters within a single auditory stream is a lack of orthogonality (Kramer, 1994b), wherein changes in one variable may influence the perception of changes in another variable. For example, numerous studies have shown that the auditory dimensions of pitch, loudness, and timbre interact perceptually (Garner, 1974; Grau & Kemler-Nelson, 1988; Kemler-Nelson, 1993; Melara & Marks, 1990a; Melara, Marks, & Potts, 1993; Neuhoff & McBeath, 1996; Neuhoff, McBeath, & Wanzie, 1999; Pitt, 1994). Changes in any of these perceptual dimensions can influence perception of changes in the others. Thus, when acoustic frequency, intensity, and spectral content are used to represent variable data in an auditory display, a distorted perception of the underlying data is a potential result.

Further complicating the effective design and use of sonification and auditory display is evidence that even stimulus changes within a single auditory dimension can be perceived differently on the basis of the direction and duration of the stimulus change. For example, in judging the magnitude of loudness change that occurs in a dynamic sound, listeners exhibit perceptual asymmetries between rising and falling intensity, despite intensity changes of equal physical magnitude (Canévet & Scharf, 1990; Neuhoff, 1998; 2001; Reinhardt-Rutland, 1995; Stecker & Hafter, 2000; Teghtsoonian, Teghtsoonian, & Canévet, 2000). Other work conducted particularly within the context of auditory display has shown that relative pitch and dynamic changes in pitch can have perceptual interactions that are similar to those that occur between different auditory dimensions such as pitch and loudness (Walker & Ehrenstein, 2000). These findings present a challenge to the design and application of effective auditory display and sonification. If a change in one variable can influence perceived changes in another, or if an increase of a given magnitude in an acoustic dimension is perceived as being different in size than a decrease of the same physical magnitude, then similar distortions in perception of the underlying data are likely to occur.

At the perceptual level, Grau and Kemler-Nelson (1988; Kemler-Nelson, 1993) argued that the dimensions of pitch, loudness, and timbre are integral and processed in a holistic manner. The suggestion is that, at least initially, listeners do not have primary access to these dimensions and cannot differentiate between changes in pitch, loudness, or timbre. Changes in any of these dimensions are perceived as overall holistic change in the entire signal. Nonetheless, the weight of evidence now appears to support the opposing view that listeners *can* access these dimensions and process auditory information in a single stream analytically (Melara & Marks, 1990a, b; Melara et al., 1993; Neuhoff et al., 1999). However, this is not to say that dimensional interaction does not take place or that the overall acoustic change has no impact on the perception of individual acoustic variables. This opposing view proposes that auditory dimensional interaction stems from the context created by one dimension (e.g., pitch) in which another dimension (e.g., loudness) is perceived. For example, a loud sound is perceived differently in the context created by high pitch than the context created by a lower pitch (Melara & Marks, 1990a, b; Neuhoff et al., 1999).

Garner (1974) proposed a set of converging operations (speeded sorting, restricted classification, and dissimilarity scaling) that have been used to determine whether a set of perceptual dimensions interacts. Participants are presented with stimuli that vary along two dimensions, such as pitch and loudness, and are instructed to attend to one dimension and ignore the other. If discrete variation of the unattended dimension influences performance on the attended dimension, the two dimensions are said to interact. Early experiments used only static stimuli and demonstrated that pitch and loudness do indeed interact perceptually (Grau & Kemler-Nelson, 1988; Kemler-Nelson, 1993; Melara & Marks, 1990a, b). So for example, if listeners are asked to attend to pitch and classify tones as being high or low, they are faster and make fewer errors on trials that have congruent pitch and loudness (high pitch and high loudness or low pitch and low loudness) than those that are incongruent (high pitch and low loudness or low pitch and high loudness). Similar results occur if listeners are asked to attend to loudness and ignore pitch.

More recent work has shown that under dynamic conditions, pitch and loudness interact in a similar manner. Changing pitch influences judgments of loudness change and changing loudness influences judgments of pitch change (Neuhoff & McBeath, 1996; Neuhoff et al., 1999). When listeners are asked to judge the magnitude of pitch change in a dynamic signal, they perceive the change to be greater when the direction of intensity change is consistent with the direction of frequency change. When frequency and intensity change in opposite directions, the magnitude of change in either dimension is judged to be less than when they change in the same direction. Furthermore, these effects can be in the opposite direction of those predicted by the traditional equal-loudness and equal-pitch contours (e.g., Robinson & Dadson, 1957; Stevens, 1934, 1935; Stevens & Davis, 1936).

The Limitations of Previous Work

Perhaps the most conspicuous oversight in the perceptual research on interacting auditory dimensions has been the failure to use stimulus tones that exhibit dynamic changes in pitch, loudness, and timbre. Although there is some recent work on the interaction of perceptual dimensions under dynamic conditions (Neuhoff & McBeath, 1996; Neuhoff et al., 1999; Walker & Ehrenstein, 2000), the vast majority of dimensional interaction research has used static stimuli. This is curious, given that almost all sounds in the natural environment are dynamic in at least some sense. Even with sounds that are ostensibly static, there can be perceived dynamic change due to motion by the source or the listener. One of the distinct advantages that the auditory system has over the visual system is in the processing of temporal information. Indeed, one of the primary conditions in which sonification is more advantageous than visualization of data is when the temporal characteristics of the task make it better suited for the ears than for the eyes (Kramer, 1994a). Thus, given that sonification and auditory display are particularly suited for the temporal representation of complex, dynamic data sets, it seems appropriate to conduct further work on the dynamic interaction of auditory dimensions as they pertain to sonification and auditory display.

A second limitation of the previous work as it pertains to auditory display is that the phenomenon of dimensional interaction has been couched almost exclusively in terms of selective attention. In almost every study of dimensional interaction to date, participants have been instructed to attend to a single dimension and ignore changes in a second "irrelevant" dimension. Although this technique is perhaps useful in investigating the perceptual foundations of dimensional interaction, it is a scenario that rarely occurs in the context of actual auditory displays and sonifications. More typically, users are required to monitor the states of two or more simultaneously changing variables. Thus, it seems reasonable to examine dimensional interaction under a paradigm in which listeners are instructed to monitor more than one auditory dimension.

Finally, Helmholtz (1866/1925) suggested that we have difficulty attending to sensations per se, although we easily perceive the objects and events in our environment that give rise to those sensations. Similarly, Gaver (1993) suggested a rubric for describing the perception of sounds according to their source attributes. Implicit in these two views is the idea that the perception of auditory objects and events at particular spatial locations takes precedence over the perception of dimensions such as pitch, loudness, and timbre. Yet, most of the experiments on the interaction of perceptual dimensions have been designed such that participants are required to attend to and report changes in perceptual dimensions per se. In the context of auditory display, not only are listeners required to monitor changes in these dimensions, but they are also required to translate these changes into the appropriate scale for the data being displayed. A pilot, for example, monitoring changes in airspeed that are represented by changes in pitch, must first attend to the changes in pitch and then translate these changes into meaningful units of velocity. This intermediate translation process is missing from almost all previous experiments on the interaction of perceptual dimensions. However, it is a critical question in applying the principles of dimensional interaction to the design of auditory display.

Other important research questions that have scarcely been addressed involve the relationships between the conceptual characteristics of the data and the perceptual characteristics of the acoustic signal, and the scaling factors that should be used in representing changes in data with changes in acoustic attributes. Walker (2000; Walker, Kramer, & Lane, 2000) outlined the following questions and presents data that begin to illuminate this issue:

> (1) What is the best sound parameter to use to represent a given data type? (2) Should an increase in the sound dimension (e.g., rising frequency) represent an increase or a decrease in the data dimension? (3) How much change in the sound dimension will represent a given change in the data dimension? (Walker, 2000, p. iii)

The Current Experiments

In the current study, we examined some of the pertinent questions regarding the interaction of auditory dimensions as they impact auditory display. In Experiment 1, we addressed the issue of selective attention. Listeners were presented with dynamic tones that changed in frequency and intensity and were asked to indicate the amount of overall change that occurred in the entire stimulus, not just one particular auditory dimension. Essentially, we asked the question: Do pitch and loudness still interact when listeners are asked to listen to overall stimulus change? In Experiment 2, we continued the examination of selective attention and introduced the issue of applying data labels to auditory dimensions. Listeners were presented with changing pitch and loudness and made estimates of real-world data values on the basis of the perceived acoustic changes. Yet, unlike the methods used in previous work, listeners in Experiment 2 were instructed to attend to changes in both pitch and loudness. In Experiment 2 we asked: What is the nature of pitch and loudness interaction when real-world labels are applied, and listeners are asked to monitor and report both dimensions? In Experiment 3, we added to our examination the effect of increased stimulus duration and complexity on the dimensional interaction of pitch and loudness. Much of the auditory dimensional interaction work has used static stimuli with durations of 100–200 ms. Some dynamic experiments have used stimuli that are up to 6 s in length. However, in these experiments the dimensional change has always been unidirectional within a given stimulus dimension. In Experiment 3, we increased the stimulus duration over static experiments by a factor of 10 and introduce a frequency modulation technique that more closely approximates the variability of real-world data values that are represented in auditory display. In Experiment 3 we asked the question: What is the nature of pitch and loudness interaction when the stimuli are longer and more realistic?

Experiment 1

Method

Participants. Twelve[1] undergraduate psychology students with normal hearing served as participants. Each received course credit for participation. None were aware of the hypothesis being tested.

Stimuli and apparatus. Stimuli were generated by a 16-bit sound card in a Pentium PC and fed directly to Sony MDR-V600 headphones. The frequency response of the headphones was 5 Hz–30 kHz. All headphone level measurements were made with a flat plate coupler with the sound meter microphone 1.27 cm from the center of the speaker element and used the A-weighted scale. Responses were made by using a standard two-button computer mouse while viewing a 38.10 cm CRT display. Listeners were presented with sounds that changed concurrently in frequency and intensity (either rising or falling) for 2.5 s. All stimuli were triangle waveforms and had a sampling rate of 44.1 kHz. Rising and falling frequency and intensity were crossed to create the four different types of sounds; two congruent change conditions (rising–rising and falling–falling), and two incongruent change conditions (rising–falling and falling–rising; see Figure 1). Rising intensity change was from 60 dB to 80 dB, falling was from 80 dB to 60 dB. Rising frequency was from 200 Hz to 240 Hz, falling was from 240 Hz to 200 Hz. All sounds were pulsed on and off, with 50-ms signal bursts interspersed with 75-ms periods of silence.

Design and procedure. Each listener heard each of the four types of sounds 10 times for a total of 40 trials, all presented in random order. The listener's task was to indicate how much each sound appeared to change overall (in pitch and loudness combined) on a visual analog scale by moving a cursor on a computer screen to indicate the amount of change they heard in each sound. Moving the cursor to the left end of the scale

---

[1] After an initial analysis with a smaller *n*, participants in each experiment were added to clarify the results.

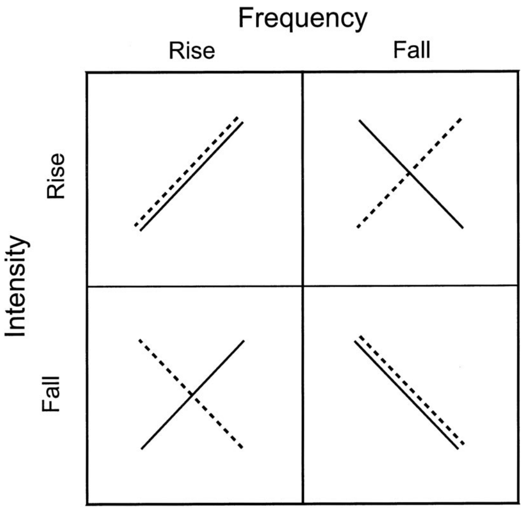

Figure 1. A schematic diagram of the four types of stimuli used in Experiments 1 and 2. Stimulus details appear in the text. Solid lines represent frequency. Hatched lines represent intensity.

indicated no change, moving the cursor all the way to the right indicated maximal change. Listeners could move the cursor to any spot between the two poles. Values on the scale were hidden from the participant but ranged from 0 to 100 for purposes of analysis. The 10 responses in each condition for each participant were averaged so that each participant contributed one data point in each condition for the analysis.

Results and Discussion

The results of Experiment 1 are shown in Figure 2. Despite an identical amount of physical change in each condition, an analysis of variance (ANOVA) revealed a main effect for intensity change, $F(1, 11) = 11.12$, $p < .01$, $f = .65$. Cohen's f is a measure of effect size where, by convention, values of at least .10, .25, and .40, correspond to small, medium, and large effect sizes, respectively (Kirk, 1995). Thus, sounds that rose in intensity appeared to change substantially more than those that fell. There was no similar main effect for frequency change, $F(1, 11) = 0.22$, $p > .05$. However, we did find a significant statistical interaction between changes in frequency and intensity, $F(1, 11) = 10.27$, $p < .01$, $f = .62$. When intensity rose, rising frequency sounds appeared to change more than falling. When intensity fell, falling frequency sounds appeared to change more than rising. In other words, there was a type of congruity effect. When pitch and loudness changed in the same direction, the sounds appeared to change more than when they changed in opposite directions (see Figure 2).

The results of Experiment 1 demonstrate two significant effects of auditory dimensional interaction that have implications for sonification and auditory display. First, it is clear that the relationship between the types of change that occurred within each auditory dimension affected the magnitude of overall auditory change, even though the amount of physical change in each condition was identical. When pitch and loudness changed in the same direction (both rising or both falling), the overall perceived change in the signal was perceived as greater than when they changed in opposite directions. Second, there was a directional asymmetry for loudness change. Sounds that increased in loudness were perceived to change more than those that decreased, despite the same amount of intensity change.

Previous work on pitch–loudness interaction using both static (Grau & Kemler-Nelson, 1988; Melara & Marks, 1990a, b) and dynamic stimuli (Neuhoff & McBeath, 1996; Neuhoff et al., 1999) has used a paradigm in which listeners are asked to attend to only one dimension. In these studies, listeners show greater speed and accuracy and judge the magnitude of the attended dimension to be greater when pitch and loudness are congruent. The results of Experiment 1 show that when listeners are asked to judge the overall change in an acoustic stimulus without requiring selective attention to any single dimension, analogous results obtain.

Listeners in Experiment 1 also heard rising-loudness sounds change significantly more than falling-loudness sounds. This finding is consistent with previous work that has shown a perceptual bias for rising intensity for sounds with relatively short durations. Listeners tend to perceive rising intensity sounds as louder overall and having a greater magnitude of change than equivalent falling intensity sounds (Neuhoff, 1998; Neuhoff, 2001; Stecker & Hafter, 2000). However, at longer durations there is evidence that listeners tend to underestimate intensity change (Canévet & Scharf, 1990; Canévet, Scharf, Schlauch, Teghtsoonian, & Teghtsoonian, 1999; Teghtsoonian et al., 2000).

The judgments of overall acoustic change in Experiment 1 are consistent with an auditory display in which multiple auditory dimensions are mapped onto a single data variable. The results, as they pertain to auditory display, suggest that congruous duplicate mapping of auditory dimensions onto a single data variable would provide greater discriminability and greater perceived change in magnitude of the variable being represented (Kramer, 1994b; Walker & Ehrenstein, 2000). This type of mapping might be particularly advantageous in situations where perceiving a change in a variable is a crucially important event. The directional asymmetry for loudness suggests that using continuous intensity change to represent increases and decreases in the value of a variable might create analogous distortions in perceived changes in the data.

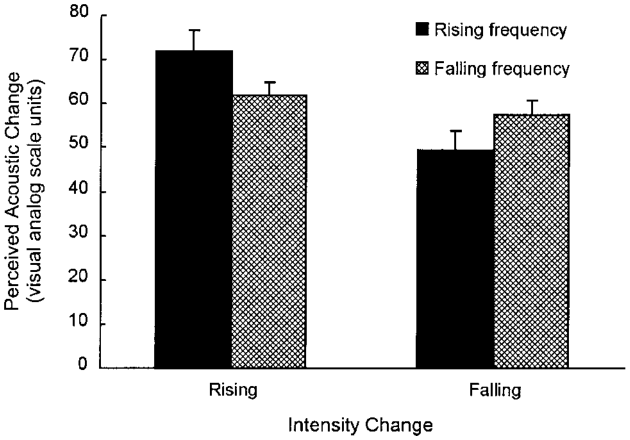

Figure 2. Average magnitude of perceived acoustic change in each condition in Experiment 1. The amount of actual physical change in each condition was identical. The statistical interaction between pitch and loudness is significant. Error bars represent one standard error.

Experiment 2

In Experiments 2 and 3, we wanted to make the listener's task more applied. We also wanted listeners to monitor and report changes in more than one auditory dimension. We did this by assigning real-world data dimensions to changes in frequency and intensity and asking listeners to report both data dimensions. Stimuli in Experiment 2 were identical to those in Experiment 1. However, this time the listener was told that changes in pitch represented changes in the price of a stock and that changes in intensity represented changes in the trading volume of that stock. Rising frequency denoted rising stock price, and rising intensity denoted rising trading volume. The task was to listen to each stimulus and make a judgment of the final stock price and trading volume on the basis of the changes in frequency and intensity.

Method

Participants. Fifteen undergraduate psychology students with normal hearing served as participants. Each received course credit for participation. None had participated in Experiment 1, and none were aware of the hypothesis being tested.

Stimuli and apparatus. Stimuli and apparatus were the same as those used in Experiment 1.

Design and procedure. Each listener heard each of the four types of sounds 10 times for a total of 40 trials, all presented in random order. Listeners were told that changes in pitch represented changes in the price of a stock and that changes in loudness represented changes in the trading volume of that stock. Rising and falling frequency denoted rising and falling stock price, and rising and falling intensity denoted rising and falling trading volume, respectively. The task was to listen to each stimulus and make a judgment of both the final stock price and trading volume on the basis of the changes in frequency and intensity that they heard. After each sound, listeners were presented with two visual analog scales similar to those used in Experiment 1. One scale was used to indicate terminal stock price. Scale values ranged from $0 to $200, and listeners were told that the initial stock price was $100 (the center of the scale). The second scale was used to indicate terminal trading volume. Scale values ranged from 0 to 200, with 10,000 shares equal to one unit on the scale. Listeners were told that the initial trading volume was 100 (the center of the scale). After indicating perceived stock price and trading volume, listeners submitted their estimates and were presented with the next stimulus.

Results and Discussion

The means for each condition in Experiment 2 are shown in Figure 3. For judgments of stock price, we found a main effect for frequency change, $F(1, 14) = 15.50$, $p < .01$, $f = .70$, indicating that listeners could perform this task and that changes in pitch did indeed guide their estimates of changes in stock price. Terminal price estimates for rising frequency trials were higher than terminal price estimates for falling frequency trials. However, we also found a main effect for intensity change that indicated an influence of trading volume on estimates of changes in stock price, $F(1, 14) = 8.10$, $p < .05$, $f = .49$ (see Figure 3). Terminal price estimates for rising intensity trials were higher than terminal price estimates for falling intensity trials. The results were consistent with the findings in Experiment 1. When frequency and intensity changed in the same direction, changes in stock price were perceived to be greater than when they changed in opposite directions. This occurred despite the same degree of frequency and intensity change in each condition. We found a significant main effect for intensity change on judgments of trading volume, $F(1, 14) = 19.00$, $p < .01$, $f = .77$. Rising intensity trials were perceived as having higher trading volume than falling intensity trials. We also found a significant effect of changes in frequency (stock price) on judgments of trading volume, $F(1, 14) = 8.25$, $p < .05$, $f = .49$. Despite an equal amount of frequency and intensity change in each condition, we found that when frequency increased, trading volume was perceived as greater than when frequency decreased.

To our knowledge, studies of dimensional interaction have not previously used methodologies in which listeners are asked to monitor perceived change in more than one perceptual dimension at a time. In almost all studies of dimensional interaction, listeners are asked to monitor a single dimension, ignore any other stimulus changes, and report or classify changes in the sensory characteristics of the dimension in question. In Experiment 2, untrained listeners successfully applied real-world data labels to these changes and accurately reported direction of simultaneous change in two dimensions of a single auditory signal. These findings suggest that, at least at an ordinal level, simultaneous monitoring of changes in auditory dimensions used to represent data is a realistic expectation of users of auditory displays. Successful performance on an ordinal scale requires only that listeners rank order stock prices and trading volumes in terms of their perceived terminal values. Differences between rising and falling conditions statistically greater than zero are sufficient for success on an ordinal scale (provided the means for rising trials are greater than the means for falling trials). On average, our listeners showed

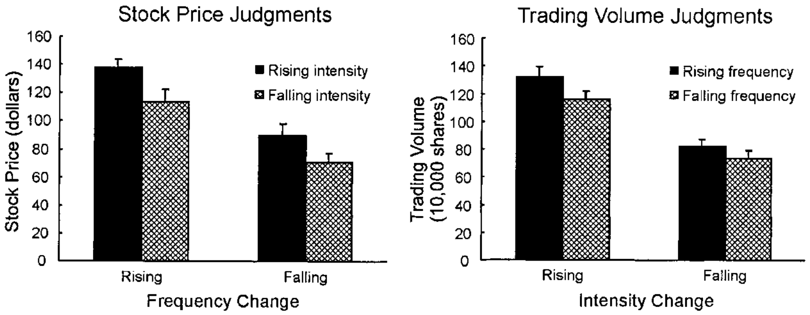

Figure 3. Mean perceived terminal stock price and trading volume in Experiment 2. Error bars represent one standard error.

accurate ordinal performance in perceiving direction of change in both frequency and intensity.

However, we also found significant distortion of perceived changes in the data that were due to the interaction of pitch and loudness. Judgments of stock price were influenced by changes in trading volume, an effect consistent with the results of Experiment 1. When price and trading volume both increased, terminal stock price was perceived to be higher than when price increased and volume decreased. When trading volume and price both fell, terminal price was perceived to be lower than when price fell and trading volume rose. This suggests that on an interval scale, where equal-interval spacing is required, or on a ratio scale where an absolute zero point is required, average performance was lacking. These findings suggest that great care should be exercised in designing auditory displays and sonifications in which the data relations are "higher-than-ordinal" (e.g., interval or ratio data). Although on average listeners could accurately detect the direction of change in each dimension, the magnitude of change was influenced by the type of change in the other dimension.

Furthermore, despite accurate ordinal performance, there may be performance costs in terms of speed and accuracy when pitch and loudness change in opposite directions. Listeners typically exhibit poorer performance when pitch and loudness are incongruent than when they are congruent (Grau & Kemler-Nelson, 1988; Melara & Marks, 1990a, b). Using dynamic stimuli, Walker and Ehrenstein (2000) have shown that similar deficits in performance occur when the interacting auditory dimensions are relative pitch and pitch change.

Despite these limitations, from the perspective of auditory display design it is encouraging that listeners could monitor directional change in two auditory dimensions at once. We specifically used untrained listeners to rule out the effects extensive experience with auditory displays. It is likely that with practice and training, the effects of dimensional interaction could be reduced. However, even highly trained musicians are not immune to the effects of auditory dimensional interaction (Pitt, 1994). Thus, it appears unlikely that practice would eliminate dimensional interaction entirely.

The duration of the stimuli in Experiments 1 and 2 was 2.5 s. Although this is much longer than stimuli in most experiments on auditory interaction, it is relatively short compared with the length of signals that occur in some auditory displays. In addition, pitch and loudness in Experiments 1 and 2 exhibited only unidirectional change within a given auditory dimension. In many auditory displays and sonifications, changes in data result in dimensional change that is much more variable. In Experiment 3, we sought to address these two issues by using stimulus tones that were longer and more variable in pitch than those used previously.

Experiment 3

Method

In Experiment 3, we used the same methodology that was used in Experiment 2. However, we extended the duration of the stimulus tones to 12 s and increased the variability of the changes in frequency. We used a frequency-modulated tone to more closely simulate the variability of real-time stock changes.

Participants. Fourteen undergraduate psychology students with normal hearing served as participants. Each received course credit for participation. None had participated in Experiments 1 or 2, and none were aware of the hypothesis being tested.

Apparatus. The apparatus was the same as that used in Experiment 1.

Stimuli. Listeners were presented with sounds that changed concurrently in frequency and intensity (either rising or falling) for 12 s. All stimuli were triangle waveforms and had a sampling rate of 44.1 kHz. Rising and falling frequency and intensity were crossed to create the four different types of sounds: two congruent change conditions (rising–rising and falling–falling), and two incongruent change conditions (rising–falling and falling–rising). Rising intensity change was from 60 dB to 80 dB, falling was from 80 dB to 60 dB. Frequency was modulated sinusoidally while rising or falling, with a center frequency that rose from 200 Hz to 300 Hz or fell from 300 Hz to 200 Hz over 12 s. The modulation depth was 50 Hz, and the rate of modulation began at 10 Hz and decreased linearly to 0 Hz as the tone progressed.

Design and procedure. Each listener heard each of the four types of sounds 10 times for a total of 40 trials, all presented in random order. Listeners were told that changes in pitch represented changes in the price of a stock and that changes in loudness represented changes in the trading volume of that stock. Rising and falling frequency denoted rising and falling stock price, and rising and falling intensity denoted rising and falling trading volume, respectively. As in Experiment 2, the task was to listen to each stimulus and make a judgment of both the final stock price and trading volume on the basis of the changes in frequency and intensity. After each sound, listeners were presented with the same two visual analog scales used in Experiment 2. One scale was used to indicate terminal stock price, and the other was used to indicate terminal trading volume. After indicating perceived stock price and trading volume, listeners submitted their estimates and were presented with the next stimulus.

Results and Discussion

The means for each condition in Experiment 3 are shown in Figure 4. Once again, we found results consistent with a dynamic interaction of pitch and loudness perception. For judgments of stock price, we found a main effect for frequency change, $F(1, 13) = 21.20$, $p < .01$, $f = .85$, indicating that changes in frequency guided estimates of changes in stock price. When frequency rose, terminal price was perceived as higher than when frequency fell. However, we also found a main effect for intensity change, $F(1, 13) = 5.33$, $p < .05$, $f = .39$, that indicated an influence of trading volume on estimates of changes in stock price. When intensity rose, terminal price was perceived as higher than when intensity fell (see Figure 4). Furthermore, the effect of trading volume on stock price was greater for rising price trials than for falling price trials, as indicated by a significant interaction between frequency and intensity change, $F(1, 13) = 16.60$, $p < .01$, $f = .75$. Judgments of trading volume were significantly influenced by changes in intensity, indicating that listeners could use intensity change to track changes in trading volume, $F(1, 13) = 26.53$, $p < .01$, $f = .95$. However, we also found a significant effect of frequency change on judgments of trading volume, $F(1, 13) = 8.30$, $p < .05$, $f = .51$. Trading volume was perceived as greater when frequency (i.e., stock price) increased than when it decreased.

The results of Experiment 3 show that concurrent frequency and intensity change within a single auditory stream can lead to distortions in perceived data change in signals that exhibit longer durations and greater variability than those typically used in dimensional interaction experiments. Thus, the effects of dimensional interaction persist even in situations that approximate those used in auditory displays and sonifications.

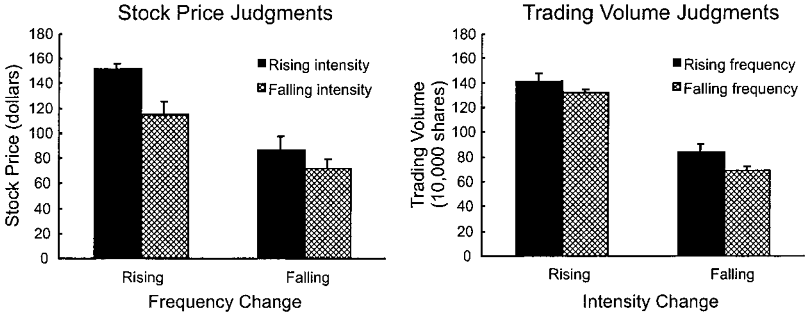

Figure 4. Mean perceived terminal stock price and trading volume in Experiment 3. Error bars represent 1 standard error.

In addition, the main findings of Experiment 2 were replicated. Listeners could track overall direction of pitch and loudness change simultaneously and could apply real-world data labels to these dimensions with accurate performance at the ordinal level of representation. This was true even when frequency changed in a highly variable manner and exhibited modulation within the signal. Conversely, the significant effect of trading volume on perceived changes in stock price and the significant interaction between the two dimensions suggests that if pitch and loudness are used to represent different data dimensions under these types of conditions, a distorted perception of actual data change is likely to occur.

General Discussion

The current results demonstrate a substantial perceptual interaction of dynamic pitch and loudness that has implications for auditory display and sonification. Congruent directional changes in pitch and loudness were perceived as greater in magnitude than equivalent incongruent directional change. Rising intensity sounds were perceived to change more than equivalent falling intensity sounds. Furthermore, the effects of dimensional interaction persisted when listeners were given tasks that more closely approximate those used in real-world auditory display situations, including monitoring more than one dimension simultaneously, assigning data labels to auditory dimensions, and monitoring directionally variable dynamic change within a dimension.

When frequency and intensity changed in the same direction in Experiment 1, the perceived amount of total stimulus change was greater than when the dimensions changed in opposite directions, despite an equal amount of physical change in all conditions. The disparity between physical and perceived change suggests that in situations where precision is critical, and accurate representation of changes in data are desired, great care should be taken to minimize the effect of dimensional interaction on perceived changes in variables. There was also a significant difference between rising and falling changes in intensity. Across the two directions of frequency change, sounds that got louder were perceived to change more than sounds that got softer. These results are consistent with previous work that shows a perceptual bias for rising intensity tones (Neuhoff, 1998; Neuhoff, 2001; Stecker & Hafter, 2000). Again, the perceptual asymmetry demonstrated with physically equivalent signals suggests that lower level acoustic dimensions such as frequency and intensity pose a potential problem for representing data with simple acoustic attributes.

In addition to the perceived asymmetries in stimulus change found in Experiment 1, there are also implications for sonification applications and the design of auditory displays. In Experiments 2 and 3, listeners were given a more realistic task wherein data variables were represented by changes in frequency and intensity. In both experiments, there was a significant influence on the dimension of interest (e.g., perceived stock price) by another simultaneously changing dimension (trading volume). Previous work on the interaction of auditory dimensions has focused almost exclusively on issues of selective attention. Listeners are asked to attend to a single perceptual dimension and ignore irrelevant change in an unattended dimension. The current experiments show that when listeners are asked to attend to two relevant dimensions, dimensional interaction and the resulting distortion in data relations also occurs.

We should note that we restricted the variation in frequency and intensity to ranges that may be smaller than those used in some auditory displays. However, previous work has shown that the interaction of pitch and loudness under dynamic conditions occurs under a wide range of frequencies, intensities, and timbres (Neuhoff & McBeath, 1996; Neuhoff et al., 1999). Furthermore, the asymmetry in dynamic loudness perception (i.e., rising intensity is perceived to change more than falling intensity) increases as stimulus intensity increases from 60 dB to 90 dB (Neuhoff, 1998).

One positive implication of the current findings is that the effect of congruent change in two auditory dimensions is more salient than incongruent change, or even acoustic change in a single dimension (Neuhoff et al., 1999). Thus, in situations where changes in the state of a variable are particularly critical, duplicate mapping of frequency and intensity onto the same variable would likely provide improved performance (Kramer, 1994b; Walker & Ehrenstein, 2000). For example, simultaneous and dimensionally congruent changes in pitch, loudness, pulsing speed, and timbre have been used in the sonification of Radionuclide Ventriculography (RVG)—a noninvasive means for diagnosing heart disease by obtaining the blood volume change of the left ventricle (Kramer, 1996). Multiple mapping of acoustic variables provided an apparently effective means of discerning healthy from unhealthy hearts.

The practical significance of the perceptual interaction between pitch and loudness was not trivial. There were large effect sizes in each experiment, and distortions of the data underlying an auditory display could have grave real-world consequences. Thus, from an applied perspective, the importance of effectively mapping data variables onto acoustic variables suggests that a greater understanding of perceptual interaction should be pursued. Such an understanding will facilitate more effective sonification techniques, particularly in situations where precision is critical. The current results suggest that listeners can discern simultaneous dimensional change in pitch and loudness and use this information to represent data on an ordinal scale. The demonstrated perceptual interaction suggests that representations of interval or ratio data in this manner would prove problematic. A potential alternative to the use of lower level perceptual dimensions such as pitch and loudness in sonification might be to use higher order acoustic characteristics or patterns. Gaver (1993) argued that listeners do not normally listen to, or even easily identify, changes in pitch, loudness, and timbre. He has suggested instead that listeners hear and easily identify acoustic events and sources in the environment. Perhaps a sonification technique that uses acoustic events or source characteristics might be more successful in avoiding the pitfalls of dimensional interaction. This type of display was posited by Kramer (1994b) in which a "virtual engine" (i.e., a theoretical sonification using an automobile engine sound as the display metaphor, replete with veridical interactions of acoustic dimensions), is suggested as a display device for a data set that is unrelated to physical engines.

Finally, although the current results are limited to dimensional interaction within the auditory domain, there is evidence to suggest that cross-modal associations and congruency effects exist between audition and other modalities (Marks, 1982). Multimodal displays are used in a variety of different environments. The applied characteristics of these associations and interactions should be explored as well, in both real and virtual environments.

References

Barrass, S., & Zehner, B. (2000). Responsive sonification of well logs. In P. Cook (Ed.), *Proceedings of the International Conference on Auditory Display* (pp. 72–80). Atlanta, GA: International Community for Auditory Display.

Bronkhorst, A. W., Veltman, J. A. H., & van Breda, L. (1996). Application of a three-dimensional auditory display in a flight task. *Human Factors, 38,* 23–33.

Canévet, G., & Scharf, B. (1990). The loudness of sounds that increase and decrease continuously in level. *Journal of the Acoustical Society of America, 85,* 2136–2142.

Canévet, G., Scharf, B., Schlauch, R. S., Teghtsoonian, M., & Teghtsoonian, R. (1999, April 22). Perception of changes in loudness. *Nature, 398,* 673.

Fitch, T., & Kramer, G. (1994). Sonifying the body electric: Superiority of an auditory over a visual display in a complex, multi-variate system. In G. Kramer (Ed.), *Auditory display: Sonification, audification, and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD)* (pp. 307–326). Reading, MA: Addison-Wesley.

Flowers, J. H., Buhman, D. C., & Turnage, K. D. (1997). Cross-modal equivalence of visual and auditory scatterplots for exploring bivariate data samples. *Human Factors, 39,* 341–351.

Flowers, J. H., & Hauer, T. A. (1995). Musical versus visual graphs: Cross-modal equivalence in perception of time series data. *Human Factors, 37,* 553–569.

Garner, W. R. (1974). *The processing of information and structure.* Potomac, MD: Erlbaum.

Gaver, W. W. (1993). What in the world do we hear? An ecological approach to auditory event perception. *Ecological Psychology, 5,* 1–29.

Gaver, W. W., Smith, R. B., & O'Shea, T. (1991). *Effective sounds in complex systems: The ARKola simulation.* Paper presented at Computer–Human Interaction, New York, NY.

Grau, J. A., & Kemler-Nelson, D. G. (1988). The distinction between integral and separable dimensions: Evidence for the integrality of pitch and loudness. *Journal of Experimental Psychology: General, 117,* 347–370.

Hayward, C. (1994). Listening to the earth sing. In G. Kramer (Ed.), *Auditory display: Sonification, audification, and auditory interfaces. Proceedings of the First International Conference on Auditory Display (ICAD)* (pp. 369–404). Reading, MA: Addison-Wesley.

Helmholtz, H. (1925). *Treatise on physiological optics* (3rd ed. Vol. III). New York: Optical Society of America. (Original work published 1866)

Hermann, T., Meinicke, P., & Ritter, H. (2000). *Principle curve sonification.* In P. Cook (Ed.), *Proceedings of the International Conference on Auditory Display* (pp. 81–86). Atlanta, GA: International Community for Auditory Display.

Kemler-Nelson, D. G. (1993). Processing integral dimensions: The whole view. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1105–1113.

Kirk, R. E. (1995). *Experimental design: Procedures for the behavioral sciences.* Pacific Grove, CA: Brooks/Cole.

Kramer, G. (1994a). An introduction to auditory display. In G. Kramer (Ed.), *Auditory display: Sonification, audification, and auditory interface, Santa Fe Institute studies in the sciences of complexity, Proceedings XVIII* (pp. 1–77). Reading, MA: Addison-Wesley.

Kramer, G. (1994b). Some organizing principles for representing data with sound. In G. Kramer (Ed.), *Auditory display: Sonification, audification, and auditory interface, Santa Fe Institute studies in the sciences of complexity, Proceedings XVIII* (pp. 185–222). Reading, MA: Addison-Wesley.

Kramer, G. (1996, November). *Mapping a single data stream to multiple auditory variables: A subjective approach to creating a compelling design.* Paper presented at the International Conference on Auditory Display, Palo Alto, CA.

Kramer, G., Walker, B., Bonebright, T., Cook, P., Flowers, J., Miner, N., et al. (1999). *Sonification report: Status of the field and research agenda* (ISBN No. 0-9670904-0-7). Palo Alto, CA: International Community for Auditory Display.

Lunney, D., & Morrison, R. (1981). High technology laboratory aids for visually handicapped chemistry students. *Journal of Chemistry Education, 58,* 228–231.

Marks, L. E. (1982). Bright sneezes and dark coughs, loud sunlight and soft moonlight. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 177–193.

McKinley, R. L., & Ericson, M. A. (1997). Flight demonstration of a 3-D auditory display. In R. H. Gilkey & T. R. Anderson (Eds.), *Binaural and spatial hearing in real and virtual environments* (pp. 683–699). Mahwah, NJ: Erlbaum.

McKinley, R. L., Ericson, M. A., & D'Angelo, W. R. (1994). 3-Dimensional auditory displays: Development, applications, and performance. *Aviation, Space, & Environmental Medicine, 65*(5, Sect. 2, Suppl), A31–A38.

Melara, R. D., & Marks, L. E. (1990a). Interaction among auditory dimensions: Timbre, pitch, and loudness. *Perception and Psychophysics, 482,* 169–178.

Melara, R. D., & Marks, L. E. (1990b). Perceptual primacy of dimensions: Support for a model of dimensional interaction. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 398–414.

Melara, R. D., Marks, L. E., & Potts, B. C. (1993). Early-holistic processing or dimensional similarity? *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1114–1120.

Neuhoff, J. G. (1998, September 10). Perceptual bias for rising tones. *Nature, 395,* 123–124.

Neuhoff, J. G. (2001). An adaptive bias in the perception of looming auditory motion. *Ecological Psychology, 13,* 87–110.

Neuhoff, J. G., & McBeath, M. K. (1996). The Doppler illusion: The influence of dynamic intensity change on perceived pitch. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 970–985.

Neuhoff, J. G., McBeath, M. K., & Wanzie, W. C. (1999). Dynamic frequency change influences loudness perception: A central, analytic process. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1050–1059.

Pereverzev, S. V., Loshak, A., Backhaus, S., Davis, J. C., & Packard, R. E. (1997, July 31). Quantum oscillations between two weakly coupled reservoirs of superfluid 3He. *Nature, 388,* 449–451.

Pitt, M. A. (1994). Perception of pitch and timbre by musically trained and untrained listeners. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 976–986.

Reinhardt-Rutland, A. H. (1995). Evidence for frequency-dependent and frequency-independent components in increasing- and decreasing-loudness aftereffects. *Journal of General Psychology, 122,* 59–68.

Robinson, D. W., & Dadson, R. S. (1957). Threshold of hearing and equal-loudness relations for pure tones, and the loudness function. *Journal of the Acoustical Society of America, 29,* 1284–1288.

Saue, S., & Fjeld, O. K. (1997, November). *A platform for audiovisual seismic interpretation.* Paper presented at the International Conference on Auditory Display, Palo Alto, CA.

Speeth, S. D. (1961). Seismometer sounds. *Journal of the Acoustical Society of America, 33,* 909–916.

Stecker, G. C., & Hafter, E. R. (2000). An effect of temporal asymmetry on loudness. *Journal of the Acoustical Society of America, 107,* 3358–3368.

Stevens, S. S. (1934). The attributes of tones. *Proceedings of the National Academy of Sciences, USA, 20,* 457–459.

Stevens, S. S. (1935). The relation of pitch to intensity. *Journal of the Acoustical Society of America, 6,* 150–154.

Stevens, S. S., & Davis, H. (1936). Psychophysiological acoustics: Pitch and loudness. *Journal of the Acoustical Society of America, 8,* 1–13.

Teghtsoonian, R., Teghtsoonian, M., & Canévet, G. (2000). The perception of waning signals: Decruitment in loudness and perceived size. *Perception & Psychophysics, 62,* 637–646.

Turnage, K. D., Bonebright, T. L., Buhman, D. C., & Flowers, J. H. (1996). The effects of task demands on the equivalence of visual and auditory representations of periodic numerical data. *Behavior Research Methods, Instruments & Computers, 28,* 270–274.

Valenzuela, M. L., Sansalone, M. J., Krumhansl, C. L., & Streett, W. B. (1997, November). *Use sound for the interpretation of impact-echo signals.* Paper presented at the International Conference on Auditory Display, Palo Alto, CA.

Walker, B. (2000). *Magnitude estimation of conceptual data dimensions for use in sonification.* Unpublished doctoral dissertation, Rice University, Houston, TX.

Walker, B. N., & Ehrenstein, A. (2000). Pitch and pitch change interact in auditory displays. *Journal of Experimental Psychology: Applied, 6,* 15–30.

Walker, B. N., Kramer, G., & Lane, D. M. (2000). *Psychophysical scaling of sonification mappings.* In P. Cook (Ed.), *Proceedings of the International Conference for Auditory Display* (pp. 99–104). Atlanta, GA: International Community for Auditory Display.

Wegner, K. (1998). *Surgical navigation system and method using audio feedback.* Paper presented at the International Conference on Auditory Display, Glasgow, Scotland.

Received February 19, 2001
Revision received August 6, 2001
Accepted August 22, 2001 ■