Trust Your Heart: Assessing Cooperation and Trust with Biosignals in Computer-Mediated Interactions

Nick Merrill
BioSENSE Lab, UC Berkeley
School of Information
Berkeley, California, USA
ffff@berkeley.edu

Coye Cheshire
BioSENSE Lab, UC Berkeley
School of Information
Berkeley, California, USA
coye@berkeley.edu

ABSTRACT
This study examines heartrate sharing in the context of a trust-building game. Through quantitative and qualitative analyses, we find that "elevated" (versus “normal”) heartrate of an exchange partner is associated with negative mood attributions and reduced cooperation in a social dilemma game. To investigate how specific our findings are to heartrate (as opposed to some other "elevated" signal collected from the body), we replicate our initial experiment with an unfamiliar biosignal, “skin reflectivity”. We find that both heartrate and the unfamiliar biosignal are associated with negative mood attributions, but we observe a decrease in cooperative behavior only with elevated heartrate. Qualitative results indicate that individuals may learn an association between our unfamiliar biosignal and the cooperative, trusting behavior of their partner. Our findings highlight the role prior beliefs can play in shaping interpretations of a biosignal, while suggesting that designers can, perhaps inadvertently, train users to associate signals with social meanings. We discuss implications for how wearable sensors can mediate social interactions.

Author Keywords
Computer-mediated communication; biosignals; heartrate; trust; cooperation

ACM Classification Keywords
H.5.3 Information Interfaces and Presentation: Group and Organization Interfaces

INTRODUCTION
As of 2016, several apps allow users to share their heartrate with their friends, leading some [19] to wonder why anyone would anyone want to do such a thing. In fact, heartrate is a potentially rich signal for designers. The meaning of a heartrate in any given context is at once socially informative [14, 25] and highly ambiguous [20]. While existing work establishes the centrality of context in shaping interpretations of a shared heartrate [20, 25], our understanding of what heartrate *can* mean as a computer-mediated cue, and how interpretations of heartrate might affect social behavior, remains extremely limited.

After all, heartrate is not just some number. The sense of one’s heartbeat is an integral feature of the human experience, and people’s associations with it range from intimacy [16] to anxiety [11] to sexual arousal [28]. Many heartrate sharing applications rely on these associations, asking users to ascribe contextual meanings to heartrate [5, 9, 23, 25], often with the aim of increasing intimacy [16]. For example, Apple Watch's heartrate sharing feature seems to be aimed at romantic partners [17]. The mobile app Cardiogr.am's website asks users, "What's your heart telling you?" [5]. These applications, along with many others, rely on the fact that people will imbue their heartrate data with emotional, and highly contextual interpretations. Given the relatively large number of wearables with embedded heartrate monitors (watches, bands, even earbuds) [27], it is unsurprising that designers are looking beyond fitness and health for ways to increase user engagement with these devices [5, 17]. However, it is not clear how individuals will interpret a shared biosignal (e.g. heartrate) in different contexts of social interaction.

In this study, we focus on uncertain social interactions, in which individuals can cooperate and develop trust. The development of cooperation and trust in the presence of uncertainty is both of theoretical importance to CSCW [3, 6, 7, 20, 24], and of practical importance, considering many technology-mediated heartrate-sharing projects aim to “enhance” cooperation, trust and intimacy [15, 25]. We apply quantitative and qualitative analyses to an iterated prisoner's dilemma game, in which heartrate information (“elevated” or “normal”) was shared between players (Study 1). In a follow-up study (Study 2), we replicate our initial study, but replace heartrate with an unfamiliar biosignal, “Skin Reflectivity Index (SRI).”

Our results raise important questions for applications that transmit sensor-derived signals socially between users. For signals with strong cultural associations, people’s prior beliefs will color their interpretations, and social outcomes may or may not be positive. In the case of novel signals, on the other hand, our results imply that designers can (perhaps

CSCW '17, February 25-March 01, 2017, Portland, OR, USA

ACM 978-1-4503-4335-0/17/03···$15.00
DOI: http://dx.doi.org/10.1145/2998181.2998286

inadvertently) *teach* users to associate these biosignals with social meanings. This effect could be viewed as beneficial, depending on design objectives. It could also be dangerous if designers suggest, perhaps even inadvertently, interpretations that lead to discrimination.

WHAT CAN A HEARTRATE MEAN?

Heartrate has deep-rooted cultural significance in many societies, and near-universal familiarity as a feature of our lived experiences. Building on associations with intimacy and love, many heartrate sharing applications have aimed to “enhance” social connectedness by fostering feelings of intimacy [16, 18, 21, 30] between people.

What heartrate means as a computer-mediated cue, however, is ambiguous, its potential interpretations varying widely in different contexts [20, 25]. Boehner et al (2007) argue for the intrinsic ambiguity of sensor data as a resource in design, particularly in systems that seek to use these data to express emotion [2]. Many technology probes corroborate this stance, relying on users to project socially contextual meanings around a transmitted heartrate [5, 25, 29]. Consequently, more recent work has challenged the notion that the social consequences of transmitting physiological data will always result in increased trust and intimacy [15, 25]. There remains little work, however, on how the potential ambiguity of a heartrate signal is resolved in social conditions of risk and uncertainty.

HEARTRATE, MOOD, COOPERATION AND TRUST

Past studies on heartrate sharing indicate that people do read social emotional and mood cues in other people’s heartrate [20, 25]. In this study, we focus on the association between heartrate information and assessments of mood, as well as behavioral outcomes of cooperation and trust towards another person. Throughout this paper, we refer to the other person that one might trust or cooperate with in social interaction as an *exchange partner*, or simply ‘partner’ for short.

In informing our hypotheses about interpersonal interactions with heartrate, we look to studies on interpretations of one’s own heartrate, which empirical research has explored more deeply [22, 28] (see [33] for a review). These studies have revealed that elevated heartrate can yield negative interpretations about one’s own mood. Generally when individuals believe that their heartrate is elevated, they often believe their mood and emotions to be more negative. Thus, we apply this same logic to how individuals will interpret the elevated heartrates of others in uncertain social interactions:

H1. Participants who see a consistently elevated heartrate from their partner will rate their partner more negatively on mood attributes, compared to participants who see a consistently normal heartrate in uncertain and risky social interactions.

If elevated heartrate has a negative connotation with mood, then elevated heartrate may increase uncertainty about the behavior of one’s partner as well. When people know that their partner has an elevated heartrate in an uncertain, risky interactions, they may take actions to protect themselves against potential losses. In trust-building situations, individuals take small risks with other people (entrustment behavior) and learn whether the other person honors that trust or not (cooperative behavior). Thus, individuals have two different ways to respond to increased uncertainty about their partners’ behavior in trust situations: 1) reduce the amount they entrust to their partners, or 2) decrease their willingness to cooperate with the partner [6, 7]. Since we expect elevated heartrate to have pre-existing connotations with negative attributes, we predict that individuals will entrust and/or cooperate less to protect themselves from potential harm when the partner has an elevated vs. a normal heartrate.

H2. Participants who see an elevated heartrate from their partner will (a) trust less, and (b) cooperate less with the partner in uncertain and risky social interactions compared to participants who see a normal heartrate.

STUDY 1: SHARING HEARTRATE IN A RISKY, UNCERTAIN INTERACTION

In order to test our hypotheses, we conducted a repeated trust experiment with shared heartrate information. Trust games present participants with financial incentives to pay attention to their partner’s decisions over time, and provide means for operationalizing trust and cooperation in the presence of uncertainty [6].

The overall design of the trust game involves anonymous pairs of fixed partners making repeated decisions to entrust valued resources to the partner, and to return (cooperate) or keep (defect) the points entrusted by the other partner. Importantly, individuals can make the highest amount of money when they entrust many points to a partner and the partner returns these points. This creates an uncertain social situation in which participants are trying to earn real money by repeatedly taking risks (entrusting points) to a partner. Since the partners are making the same decisions to entrust and keep/return points from the other partner, these are mutually-dependent social interactions.

Experimental Design and Methods

We operationalized an uncertain social interaction situation using a trust game called the Prisoner’s Dilemma with Dependence (PDD) [6, 7]. The PDD game allows individuals to control the amount of risk that they want to take with their partner by choosing how many points to entrust, followed by a second decision to either keep or return whatever has been entrusted by their partner. Thus, the PDD game separates trust behavior (choosing how much to entrust to a partner) from cooperative behavior (choosing to return or keep what a partner entrusted).

In each round of the PDD game, participants were given an initial endowment of 10 points. Each participant decided whether to entrust any number of points to their partner,

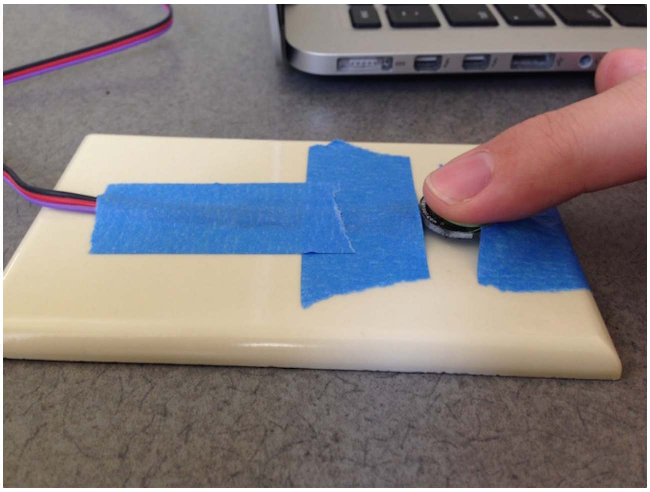

Figure 1. The heartrate monitor. Participants were told to place their finger on the monitor to take a reading while viewing their partner's decisions during the previous turn.

from zero to ten. Then, participants found out at the same time whether their partner had entrusted them with any of their own points, and if so, how many. Next, each participant decided whether to keep the points entrusted to them (defection) or return them (cooperation). The participants could not return only a portion of the entrusted points, only all or none of them. If the points were returned to the partner, they were automatically doubled in value for that participant.

After all participants made decisions about returning or keeping any points that had been entrusted to them, they were then asked to place their finger on the heartrate monitor for a few seconds in order to get a pulse reading (Figure 1). Participants then viewed the summary of point calculations for the round. Subsequently, participants viewed a visual display of the partners' recent heartrate (Figure 2). The final point calculation for the round included any of the initial allotment of points remaining after the trust decision, plus and points that the participant kept from their partner if they decided not to return them. In addition, players received points for any entrusted points that their partner returned, which doubled in value.

When participants arrived at the laboratory, they were given a consent form that described the nature of the study, as well as the human subjects' approval information from our university. We wanted participants to believe that they would be interacting with other real people, and this perception was enhanced by having 12-16 participants at separate computer terminals in the same large room during each experimental session. In fact, we controlled the trust and cooperation behavior of the "partner" for every participant using a simulated computer actor. As a result, no one in the study interacted with a human partner.

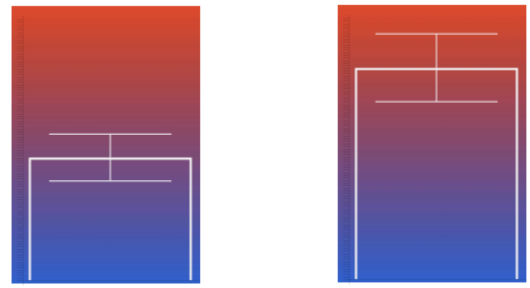

Figure 2. The heartrate visualization. After viewing the results of the previous round, participants saw a graph of what they believed to be their partner's heartrate, either normal (left) or elevated (right). Error bars fluctuated within pre-set bounds.

The simulated actor was programmed to always begin by entrusting one point on the first round, then randomly entrust up to one point above or below whatever the partner entrusted on the previous round. In addition, the simulated actor was programmed to always cooperate (i.e., return the points that were entrusted by the partner). Following [6], we chose to use a highly cooperative interaction partner in order to minimize any other forms of uncertainty in the interaction. A highly-cooperation partner does not introduce any defection behaviors that might otherwise reduce cooperation or trust from the participant (thereby hindering our ability to detect main effects from the experimental manipulation). Thus, the simulated actor was designed to reciprocate the entrusting behavior of the human participant on each round, and always cooperate no matter what the human participant chose to do.

The participants completed 20 rounds of the PDD game, but they did not know how many rounds they would play in order to eliminate end-game effects. After all rounds of the PDD game were completed, participants answered a short post-questionnaire in order to assess their attitudes and beliefs about their partner. This questionnaire included 7-point Likert-style response questions (1 = strongly disagree, 7 = strongly agree) about the partners' beliefs about the partners' anxiety (e.g., "my partner is anxious" and "my partner is calm").

As a manipulation check on the perceptions of the simulated actor's behavior, we also asked questions about the partners' game behavior ("my partner is trustworthy" and "my partner is cooperative"). Finally, we supplemented our quantitative measures with two open-ended questions: "How would you describe your partner?" and "What, if anything, did heartrate tell you about your partner during this experiment?" Participants were paid between $15-30 based on their point earnings during the game. The entire study lasted one hour.

At the end of the study, participants were debriefed on the true nature and intent of the experiment. An experimenter was available at the end of the study in case of any

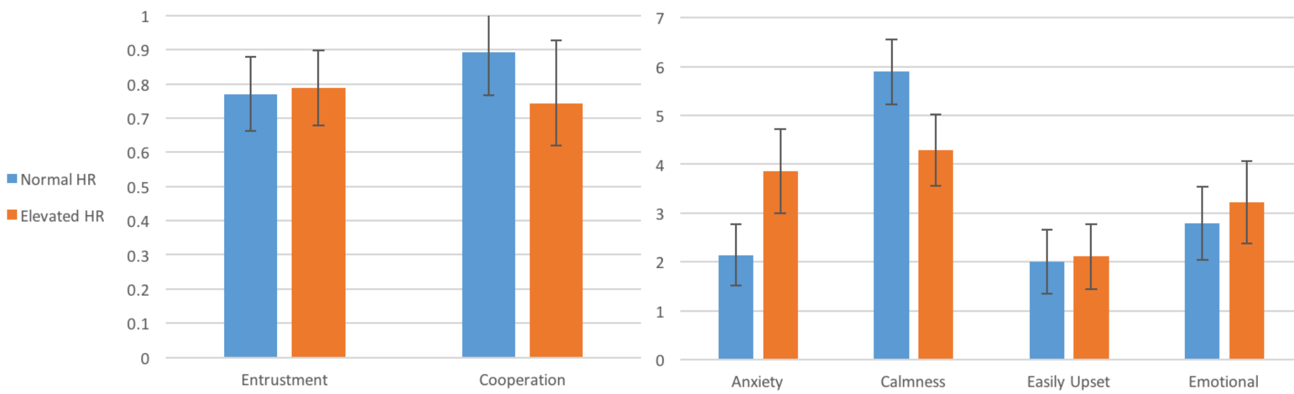

Figure 3. Means of entrustment and cooperation (left) and mood attributions (right) in elevated and normal heartrate conditions.

questions, and we provided participants with the researchers' email addresses on both the signed informed consent form, as well as the debrief form, so that they could contact us regarding any aspect of the study. We did not receive any emails or concerns from participants.

Experimental Manipulation

To assess the effect of interacting with a partner who has an elevated heartrate versus interacting with a partner who has a normal heartrate, we controlled the heartrate information that participants saw after each round of the experiment. This created a two-condition design: always normal heartrate (NH) and always elevated heartrate (EH).

Participants and Procedure

Our sample was undergraduate students recruited from the population of a large west coast public university in the United States. We contacted potential participants via email from a voluntary experimental subject pool. All participants expected to be contacted to participate in a social research study at some point during the semester, and knew that they would earn between $15-30 during this one-hour study, depending on their choices during the experiment. Fifty-six participants (56) completed the experiment, 41 women, 14 men, and one self-identified as other. The mean age of participants was 21.

Upon arrival at the laboratory, participants were guided to an individual desk with privacy walls. After signing an informed consent form, participants read written instructions on the computer which explained that they will have the opportunity to interact with a single partner for many rounds in order to examine decision making in social situations. Participants were also told that we would collect pulse (heart rate) information at designated times during the study using a simple pulse monitor that was connected to the laptop computer.

Validity Check of the Visualization

Our study aims to understand the effect of "elevated," as compared to "normal," heartrate. As such, we needed to show participants a visualization that afforded only a *relative* value for heartrate, not an exact figure (since different people may have different ideas of what number value constitutes a normal or elevated heartrate).

We designed a visualization to display a relative heartrate (Figure 2) and performed a small (n=25) face validity check to ensure that our visualization would work as intended in the actual experiment. In our short validity survey, we included three versions of the visualization, representing a mix of elevated, low and normal heartrate, and two Likert-scale questions: "The precise meaning of this graphic is ambiguous," and "I can interpret the difference between 'low', 'normal', and 'high' heartrate from this graphic," which participants answered from "Strongly Agree" to "Strongly Disagree" on a 5-point scale. We also included two open-ended questions, "Please explain what the picture is telling you about one's heartrate," and "Please explain what this picture does not tell you about one's heartrate."

We distributed this survey over an email list to students and alumni of a public, West Coast US university, and received 25 valid responses. The answers to both Likert questions indicated agreement that the visualization was both ambiguous (mean = 3.58, S.D. = 1.28) and also easily interpretable (mean = 3.41, S.D. = 1.35). Importantly, open-ended qualitative responses confirmed that the heartrate was easily understandable, but that the precise value of heartrate was ambiguous.

STUDY 1: QUANTITATIVE RESULTS

Our first hypothesis predicts that, when individuals believe that their partner has a consistently elevated heartrate, compared to a normal heartrate, they will rate the partner more negatively on mood attributes. Consistent with prior research, we found an overall strong, statistically significant effect and medium practical association between attributions and experimental condition, $F(4, 51) = 6.7$, $p < .0001$; Wilk's lambda = .66, partial eta squared =.34. Turning to the individual outcomes, we find that perceptions of the partners' anxiety is significantly higher in the EH condition ($M = 3.86$, $SD = 1.72$) compared to the NH condition ($M = 2.14$, $SD = 1.27$), $F(1, 54) = 18$, $p < .001$; partial eta squared = .25. Furthermore, participants rated their partners as significantly more calm in the NH condition ($M = 5.9$, $SD = 1.3$) compared to the EH condition ($M = 4.29$, $SD = 1.46$), $F(1, 54) = 18.71$ $p < .001$; partial eta squared =.26. On the other hand, we found no statistically significant differences for perception that the partner is “easily upset” or that the partner is “emotional” (p = n.s.). In sum, we find strong statistical and practical differences in perceptions of both anxiety and calmness, but no statistical or practical differences in perceptions of how emotional or easily upset the partner is in the two experimental conditions. Given the significant omnibus test and significant results on two of the four individual outcomes, Hypothesis 1 is partially supported.

Our second set of hypotheses predict that participants in the elevated heartrate (EH) condition will exhibit lower trusting (H2a) and/or cooperative (H2b) behavior compared to those in the normal heartrate (NH) condition. The average points entrusted by participants in the EH condition ($M = 7.88$, $SD = 2.18$) was not significantly different than the NH condition ($M = 7.7$, $SD = 2.18$), t =.28, p=n.s, one-tailed test. Thus, individuals entrusted points to their partners at approximately the same level in both conditions (Figure 3). Hypothesis 2a is not supported.

However, we found that the average cooperation rate in the EH condition ($M = .74$, $SD = .37$) was statistically significantly lower than the NH condition ($M = .89$, $SD = .25$), $t = 1.76$, $p < .05$, one-tailed test. Importantly, this result shows a medium practical effect size (Cohen’s $d = .47$), indicating a meaningful real world difference. On average, those in the normal heartrate condition cooperated 20% more than those in the elevated heartrate condition (Figure 3). Hypothesis 2b is supported.

Manipulation Checks

Since we designed the simulated actors in both conditions with trusting and always-cooperative behavior, we did not expect participants to rate the simulated actors differently in terms of the focal behaviors of cooperativeness and trustworthiness between experimental conditions. This is a critical manipulation check, since we need to rule out any perceived effect of the simulated partners' behavior in order to establish that the primary treatment (heartrate of partner) had an effect on the human participants' behavior. The omnibus test of difference in perceptions of the trustworthiness and cooperative behavior between conditions was not significant, $F(2, 53) = .21$, p = n.s.; Wilk’s lambda = .99, partial eta squared =.01. Thus, as we would expect, individuals did not indicate significant behavioral differences for the trusting, cooperative simulated actor (which was programmed to behave exactly the same in both conditions).

STUDY 1: QUALITATIVE RESULTS

At the end of our questionnaire, before the demographic questions and the debriefing, participants were presented with two open-ended questions. The first asked participants to “Tell us how you would describe your partner.” The second asked participants “What, if anything, did heartrate tell you about your partner during this experiment?” This section discusses and unpacks some of the responses that these questions elicited.

Elevated Heartrate

Many people who referred to elevated heartrate in their responses mentioned that it signaled anxiety. In some cases, participants even reflected on a negative relationship between elevated heartrate, anxiety and trust:

how excited he/she is, whether he/she cheated

It was elevated all the time so I think s/he was anxious [...] so I guess s/he did not completely trust me

These quotes further support our first hypothesis, as well as findings of past work showing that elevated heartrate typically signals anxiety and mood. In other words, elevated heartrate (and heartrate in general) seemed to be about the partner's current disposition, rather than who the partner was as a person. While the majority of those who mentioned elevated heartrate implied a causal relationship between the signal and the game context, a few did not:

My partner's heart rate was elevated the whole time, most students are stressed so that might be why.

They may have been nervous because of doing the experiment itself.

The relative rarity of skepticism about the relationship between heartrate and specific game events highlights the crucial role of framing and salience in turning what might be a disembodied signal (heartrate data) into a relevant, contextual clue. We also noted diversity in beliefs about the meaning of heartrate itself. Where almost all participants who mentioned heartrate associated it with anxiety, at least one participant had an entirely different take on his/her partner's consistently elevated heartrate:

My partner's heart rate does not change too much which indicates that he or she is very nice.

These quotes highlight overall diversity in what an elevated heartrate is capable of meaning. Even within our relatively small, and relatively homogenous sample of university students, our quotes imply a mostly negative association with elevated heartrate, but also a potentially long tail of diverse beliefs about elevated heartrate.

Normal Heartrate

Many participants said that normal heartrate indicated that the partner was "calm," "chilled out," or "not anxious."

[HR signaled] that my partner was always calm. The heart rate never fluctuated, it didn't make a difference.

They remained calm

I think it showed that my partner wasn't too nervous to see if he/she was returned the points or not, maybe because it was just an experiment or maybe because he/she wasn't worried about what result he/she was about to see was.

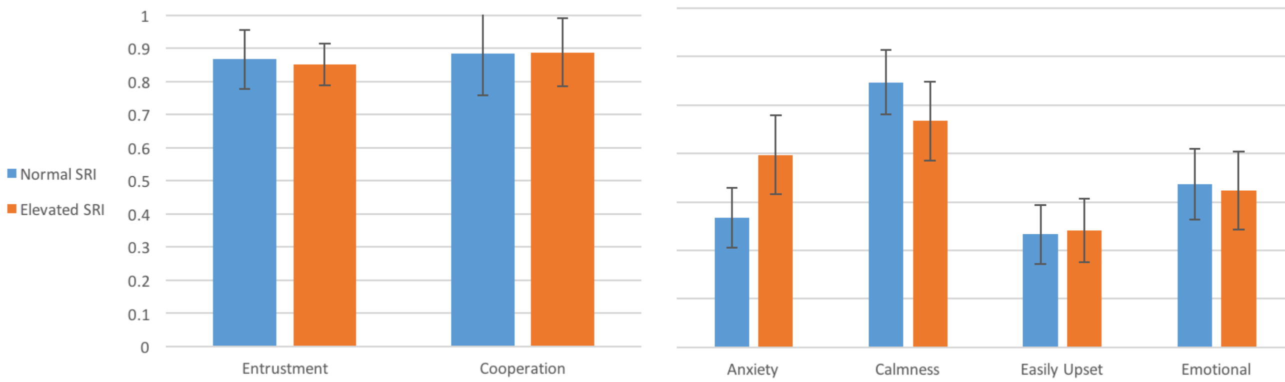

Figure 4. Means of entrustment and cooperation (left) and mood attributions (right) in elevated and normal SRI conditions.

These quotes show subjects inferring a direct connection between the heartrate signal and the attribution of a calm mood. One participant specifically mentioned that consistency of normal heartrate made their partner seem more trustworthy:

My partner's heart rate has been consistently normal throughout the experiment, so I guess s/he has no intention to cheat.

Another participant, presumably a cooperative one, thought that their partner's heartrate would have risen if s/he had not cooperated:

I think it remained the same [normal] because I paralleled my partner's actions whereas if I had contradicted them, their heartrate probably would have changed in response.

In all of the above quotes (and the vast majority of responses), participants inferred a relationship between normal heartrate and calmness. However, a few participants did not infer any relationships between behavior, moods and the signal they saw.

Heartrate did not tell me anything. My partner was average each time. I also am sure I have an elevated heart rate due to coffee consumption so I did not take my partners into consideration.

I based my decisions on their previous actions.

Not every participant explicitly inferred a calm mood from the normal heartrate signal, but most did. Taken alongside our quantitative results, our qualitative results provide evidence that subjects have used the emotional attributions they made based on their partner's normal heartrate to guide their behavior in the trust game.

STUDY 2: SHARING AN UNKNOWN SIGNAL IN A RISKY, UNCERTAIN INTERACTION

In study 1, we found that participants cooperate less with partners who have elevated heartrates in the repeated trust game, compared to those with normal heartrates. While this result supports one of our key hypotheses, it also begs another question: Is the effect we observe due to heartrate specifically, or might any elevated biosignal show the same results for negative perceptions of mood and reduced cooperative behavior towards the partner?

In our second experiment, we attempt to tease out the effect of the heartrate signal itself, compared to any "elevated" (versus "normal") signal collected from the body. We replicate the first study, except that we tell participants that our monitor device measures SRI (Skin Reflectivity Index). SRI is an unfamiliar biosignal, for which individuals should not have any prior cultural or social beliefs.

Hypotheses

Without any context for what SRI means as a signal, participants may assume that *any* biological signal that is "elevated" from normal will be negatively associated with one's mood. If this is the case, then we should observe the same general pattern of negative mood attributions and less cooperative behavior when the partner has an elevated SRI as we observed with heartrate.

On the other hand, perhaps heartrate is special due to its common social associations with mood, anxiety, and even deception. If heartrate is distinctive in this regard, then we would not observe the same significant differences between normal and elevated SRI and mood attributes, trust, and cooperation rates with the partner.

To test the effect of our unfamiliar biosignal on behavior in risky, uncertain interactions, we evaluate the exact same hypotheses from study 1 again in the context of SRI:

H3. Participants who see a consistently elevated SRI from their partner will rate their partner more negatively on mood attributes, compared to participants who see a consistently normal SRI in uncertain and risky social interactions.

H4. Participants who see an elevated SRI will have lower (a) trust rates (b) cooperation rates in uncertain and risky social interactions compared to participants who see a normal SRI.

Experimental Design and Methods

The second study was identical to the heartrate study in every way, except that we told participants we were measuring "Skin Reflectivity Index," instead of heartrate. All mentions of the word "heartrate" in our original experiment software were replaced with "SRI" and/or "Skin Reflectivity Index". We purposely did not define or explain what the SRI signal is, or what its measurements mean. All participants were debriefed on the true nature of the experiment at the conclusion of the study. This debriefing included the fact that the partner was based on idealized behavior, and "SRI" was actually just a term for heartrate, as collected by a standard light-based pulse sensor. As with the first study, participants had the ability to ask the experimenter questions at the end of the study, or send an email if they had additional questions or concerns. We did not receive any follow-up concerns from participants.

The only other variation from the first experiment is that, in the SRI experiment, we told participants to place their palms an inch above the light sensor rather than to place their fingers on the monitor. Since placing a finger on a light sensor is a familiar of measuring heartrate, this was done to reduce the possibility that participants would think that SRI is actually heartrate.

Participants

We recruited our sample for the second study from the same population and using the same method as described in study 1. Our recruitment procedures ensured that no one who participated in the first study could be recruited for the second study. Sixty-three participants (63) completed the second experiment, 40 women, 22 men, and one self-identified as 'other'. The mean age of participants was 21. Importantly, the gender distribution and age of the sample was equivalent to the first study.

STUDY 2: QUANTITATIVE RESULTS

H3 predicts that when individuals believe that their partner has a consistently elevated SRI, compared to a normal SRI, they will rate the partner more negatively on mood attributes. As with the first study on heartrate, we found an overall strong, statistically significant effect and medium practical association between attributions and experimental condition, $F(4, 59) = 4, p < .01$; Wilk's lambda = .79, partial eta squared =.21. For the individual outcomes, we find that perceptions of the partners' anxiety is significantly higher in the elevated SRI condition ($M = 3.97, SD = 1.62$) compared to the normal SRI condition ($M = 2.67, SD = 1.24$), $F(1, 62) = 12.8, p < .001$; partial eta squared = .17. Furthermore, participants rated their partners as significantly more calm in the normal SRI condition ($M = 5.5, SD = 1.3$) compared to the elevated SRI condition ($M = 4.68, SD = 1.63$), $F(1, 62) = 4.4$ $p < .05$; partial eta squared =.07. Just as with the heartrate study, we found no statistically significant differences for perception that the partner is 'easily upset' or that the partner is 'emotional' (p = n.s.). In sum, we find strong statistical and practical differences in perceptions of both anxiety and calmness, but no statistical or practical differences in how emotional or easily upset one perceives the partner to be in SRI conditions. Given the significant omnibus test and significant results on two of the 4 individual outcomes, Hypothesis 3 is partially supported.

Our final hypotheses predict that participants in the elevated SRI condition will exhibit lower trusting (H4a) and cooperative (H4b) behavior compared to those in the normal SRI condition. The average points entrusted by participants in the elevated SRI condition ($M = 8.5, SD = 1.27$) was not significantly different than the normal SRI condition ($M = 8.7, SD = 1.77$), $t = .39$, p = n.s, one-tailed test. Thus, individuals entrusted points to their partners at approximately the same level in both conditions (Figure 4). Unlike the heartrate study, however, we found no significant difference in cooperation rate between in the elevated SRI ($M = .89, SD = .21$) and the normal SRI condition ($M = .88, SD = .25$), $t = .09$, p = n.s., one-tailed test. H4a and H4b are not supported.

Manipulation Checks

As with the first study, the simulated actors in study 2 were programmed to be consistently trusting and cooperative in the elevated and normal SRI conditions. Thus, we do not expect participants to rate the simulated actors differently in terms cooperativeness and trustworthiness between experimental conditions. As expected, the omnibus test of difference in perceptions of the trustworthiness and cooperative behavior between conditions was not significant, $F(2, 61) = 3$, p = n.s.; Wilk's lambda = .91, partial eta squared =.09.

STUDY 2: QUALITATIVE RESULTS

As in the heartrate condition, participants in the SRI condition were asked open-ended questions at the end of the post-experiment questionnaire, before the demographic questions and debrief. As in the heartrate condition, participants were asked how they would describe their partner. However, unlike in the heartrate condition, participants were asked, "Recall what we were measuring with the sensor. Please describe it below." After completing this question, participants proceeded were given two more open-ended items: "What, if anything, did SRI (skin reflectivity) tell you about your partner during this experiment?" and, "As a signal, what do you believe that SRI says about another person?"

The Meaning of an Unfamiliar Biosignal

We purposely did not explain what SRI might mean in this study. Nevertheless, when asked what was being measured in SRI, some participants gave us thorough explanations:

The "reflectivity" part of SRI leads me to believe that the device is measuring how much light is reflected by a person's palms, which leads me to assume that SRI is increased when a person's hands are sweatier, and thus more covered in water, which reflects light better than simply someone's skin.

While explanations like this one indicate that participants believed our signal was real, reports of what participants thought SRI meant in the context of the game are more relevant to our analysis here. Like in the elevated heartrate conditions, and elevated SRIs were associated with either nervousness or excitement.

If the SRI reads high, it may indicate that the person expects to be betrayed in some way or is hopeful of a positive result.

I forgot what SRI stands for again. Since his/her SRI is always elevated, I would assume he/she is nervous/excited or just it's hot in here.

SRI may give insight as to how nervous or excited someone's response is to something that happens. Maybe someone with a larger range in SRI is more emotional.

These assessments of SRI are quite similar to interpretations from the elevated heartrate, and corroborate our quantitative findings that those who saw elevated SRI rate their partners as more nervous. However, the fact that these emotional assessments were similar in both elevated heartrate and elevated SRI conditions, but behavioral outcomes were different, challenges our notion that negative emotional cues caused these behavioral outcomes—a point we address in more detail in the discussion below. As in the heartrate conditions, some participants responded that SRI told them little or nothing of interest about their partner:

Nothing at all about the person other than an arbitrary value of a sensor.

Since the SRI seemed to be bouncing around in the blue range but never got into the red range (which I assume would be "abnormal" since the blue range was normal) I don't think SRI is an accurate measurement of much.

As with heartrate, people cannot always be convinced that a biosignal is informative, even after many rounds of conditioning and a highly suggestive context. However, as in the heartrate condition, responses indicating that SRI had no meaning were a clear minority in our sample.

Elevated SRI

To help explain why elevated heartrate had a chilling effect on cooperative behavior, where elevated SRI did not, we delve into the responses of participants in the elevated SRI condition. When asked what SRI told them about their partner, participants often reported nervousness or anxiety, just as we noted in the quantitative results:

[SRI shows] stress or heightened anxiety

how reactive they are, or how close to the surface their emotions are.

The nervousness of a person.

However, we noted that a significant number of participants in this condition mentioned that elevated SRI had some sort of *positive* association with behavior—even though it is also interpreted as indicating anxiety.

Elevated means they feel safe and trustful. Lower than average means they are defensive and scared.

This interpretation stands in stark contrast to elevated heartrate, which also signaled anxiety, but had a *negative* association with behavior. In explaining why participants found elevated SRI to signal cooperativeness and trust, we look toward the responses of participants who seemed to *learn* a meaning for this signal:

Well, since their SRI was always high and they always gave the money back to me, (based on these only two bits of info I know) I assume the two are correlated and an elevated SRI means that they're going to give the money back. [...] I guess it means that they're trustworthy and will do the right thing by their partner.

I cannot tell [what SRI means], but my partner's was extremely elevated for the whole experiment and s/he was good at conducting mutually beneficial transactions.

These quotes strongly suggest that, unlike for heartrate, SRI participants picked up on a pattern between their partner's always-cooperative behavior and the elevated biosignal that we displayed to them, thus filling in the gaps about what SRI meant in this context. In contrast, we found no evidence that elevated heartrate participants learned such an association in the first study, despite the fact that every participant interacted with a perfectly cooperative partner in all conditions and studies.

Normal SRI

As with those in the elevated SRI condition, many participants in the normal SRI condition identified some relationship between SRI and the other person's mood.

I think this helps identify how people are feeling internally when making decisions.

his/her mood at that point of time

[SRI shows] stress or heightened anxiety

how anxious they are.

I think our anxiety is being measured.

How anxious/nervous someone is, if their SRI is high

In some cases, participants in the normal SRI condition inferred that *elevated* SRI might have a negative meaning:

not to sure, high sri may indicate panic/fear or anger low sri may indicate calmness and contentness.

A person is less likely to trust other people if he or she has a high SRI.

Overall, the responses for both SRI conditions support the interpretation that participants learned an association between cooperative, trustworthy behavior from the partner and SRI. As we argue in the following discussion, such associations are more likely in the SRI conditions because, unlike for heartrate, participants should have no pre-existing beliefs or associations with SRI.

LIMITATIONS

Controlled, laboratory studies always come with clear advantages (such as high internal validity) and disadvantages (such as reduced external and ecological validity). Our study did not attempt to emulate a real-world interaction context with a biometric sharing device, though this is a clear next step, now that we know there are important differences in how biosignals are interpreted. Furthermore, our use of highly cooperative, computer-controlled interaction partners with stable biosignals (always high or always normal), prevents us from being able to speak to the effects of more dynamic behaviors and/or changes in biosignals over longer periods of time.

From these experiments, we also do not know how these results will transfer to other contexts, and other types of social interactions. Also, our study by nature focused on first-time, iterated interactions, both with an interface and with another unknown person. We do not know how these results might apply over the course of more personal relationships, or after repeated experiences with a specific interface in a biosignal sharing device. In addition, this research was conducted on young adults at a large public university, which is an important limitation when considering whether these results would hold across age groups and other key sources of sociodemographic variation in the larger population.

DISCUSSION

We found that both heartrate and SRI signaled negative mood to participants, including anxiety and lack of calmness. It is possible that almost any "elevated" biosignals could be associated with negative mood attributions such as anxiety and lack of calmness: many elevated signals (pulse, temperature, blood pressure) carry associations with being angry, sick, hot-headed, and a host of other negative attributions. People may default to such attributions when seeing an unknown signal that comes from the body.

Elevated heartrate had a chilling effect on cooperation, where an unfamiliar biosignal, SRI, did not. So, why did the negative mood attributions in the elevated SRI condition *not* translate into reduced cooperation, as they did for elevated heartrate?

Our results shed light on two relevant phenomena that may address this question. First, pre-existing beliefs about heartrate are powerful: even when playing with a very cooperative, trusting game partner, negative connotations surrounding elevated heartrate appear to lead individuals to cooperate less. Our results suggest that participants bring to uncertain social interactions their own expectations about what elevated heartrate means, and that these biases cannot be quickly overridden, even when behavioral evidence sends a positive message (e.g., high cooperation and trust from the partner).

Second, we find evidence that participants can "learn" a social meaning for a previously unknown signal. Our qualitative data suggest that participants in the SRI condition associated whichever signal they saw (elevated *or* normal) with cooperativeness, and trustworthiness. Unlike with heartrate, people did not have preconceived notions of how SRI should affect the social behavior of the partner, since SRI does not exist. Instead, we observe participants discovering "what SRI means" by watching their partner's behavior in relation to the biosignal. In the absence of guidelines for interpreting what SRI is or what it measures, individuals appear to fill in the gaps with available behavioral information.

If people can learn social meanings for previously unknown signals, perhaps even pre-existing connotations for familiar biosignals could change over time. After all, the meanings of a signal like heartrate are the product of associations that have been shared and developed over centuries. However, technology allows for new expressions of these ancient signals [25]. If social heartrate information became an easily accessible biosignal in trust-based interactions like negotiations, we might find its social meaning could evolve further. Unfortunately, short-term laboratory studies such as this one are unlikely to trigger or detect enduring shifts in the social meanings of familiar biosignals. We need both longer-term experiments, and mixed-methods research that can draw from rich qualitative data as well as statistically and practically significant changes in interpretations over time.

Broadly, our results raise questions about how and why unfamiliar signals take on social meanings in different contexts of interaction. Researchers in CSCW and HCI have long noted our tendency to read into cues and signals in technology-mediated communications [12]. From impact factors and citation counts in scholarly work [8, 13], to societal indices [31], to health metrics such as the body-mass index (BMI) [4], humans have a tendency to impart "real" meanings onto metrics, scales and signals – meanings that may not align with the concepts their designers aimed to measure. It is critical that we continue to question how biosignal data could shape our interpersonal interactions, and whether the outcomes will always translate into *meaningful* social information.

IMPLICATIONS FOR DESIGN AND FUTURE WORK

From research projects like the sociometer, which produce "social metrics" [32], to consumer devices like the Spire, which compute "calmness" or "focus" quotients [26], developers are throwing different biometric signals at people faster than they can learn what the signals mean in context. In the absence of strong cultural beliefs about the signal, people could produce correlative assumptions similar to the ones we observed in our experiment. Designers should take care to establish what the signals in the applications mean, or could mean. Testing the limits of what people are willing, or able, to believe, and whether these beliefs transfer between different contexts, could have wide-reaching implications for those who design interactions with wearable biosensors.

On the other hand, many research and commercial projects use signals that people might associate with commonly understood experiences (e.g. a racing heart, a sweaty palm). Designers should strongly consider how these embodied experiences might color the conclusions that users jump to, and bound what users are willing to believe.

Aside from heartrate, we do not know what many other biosignals might be associated with moods and behaviors. Similar studies with signals from, e.g., the brain [1] are a clear direction for future work. Especially interesting cases are signals for which precise or empirical meanings are still being hotly debated, such as EEG (brainwaves).

Future work should also investigate the ranges of beliefs observed among and between different populations. Even within heartrate, we find a number of different, sometimes conflicting interpretations in our qualitative findings. It is imperative that future works explore the diversity of meanings found among groups with differing expertise, or different cultural beliefs about the body. This work could help forestall issues when, e.g., internationalizing wearable applications, and could help surface beliefs about the relationship between the mind and body relevant or generative to designers of wearable applications.

Our experimental conditions involved a stable biosignal (always elevated or normal) and stable behavior (always cooperative) from the computer-controlled partner. Controlling these variables benefits the internal validity of our study, but reduces our ability to speak to whether people can mine patterns in more complex behaviors from biosignals. For example, we do not know if individuals will associate an elevated biosignal with more interactive behaviors such as the tit-for-tat strategy (i.e., do what the other person did on the last turn). Future work should examine how familiar and unfamiliar biosignals might be associated with a wider variation of complex behaviors in trust-building situations.

We also hope that researchers will investigate settings in which biosignals vary over longer time periods, perhaps with a more naturalistic technology probe study. Such a study could help us understand how prior beliefs about signals both affect and are affected by social interactions in the course of everyday life.

In general, wearable sensors can enable social interactions in which we share *more* information than is normally possible face-to-face. The ability to surface signals that are normally socially invisible (e.g. heartrate, or galvanic skin response) presents new territory for designers of computer-mediated interactions. While recent work has explored how these novel signals fit into our existing understanding of social cues [15], much work remains. Of particular relevance to CSCW and HCI researchers, future work could advance media richness theory [10] by further exploring how biological signals can inform our interpersonal appraisals with the aid of new technologies.

CONCLUSION
Our results imply that interfaces can "teach" the meaning of some biosignals, where others carry strong, pre-existing connotations that even repeated interactions cannot easily alter. In general, prior beliefs about the body (drawn from culture, lived experience) seem to shape what a biosignal can mean in a given context. However, in the absence of prior beliefs, there exists an opportunity—and a potential danger—that designers of biosignal-sharing systems can condition participants to *learn* (potentially arbitrary) associations between biosignals and social behaviors.

ACKNOWLEDGMENTS
This work was supported by funding from the National Science Foundation (ACI-VOSS 1322270). Many thanks to Noura Howell and reviewers from recent CHI and CSCW conferences for their comments on earlier versions of this research.

REFERENCES
[1] Ali, S.S., Lifshitz, M. and Raz, A. 2014. Empirical neuroenchantment: from reading minds to thinking critically. *Frontiers in Human Neuroscience*. 8, (May 2014).

[2] Boehner, K., DePaula, R., Dourish, P. and Sengers, P. 2007. How emotion is made and measured. *International Journal of Human-Computer Studies*. 65, 4 (2007), 275–291.

[3] Bos, N., Olson, J., Gergle, D., Olson, G. and Wright, Z. 2002. Effects of four computer-mediated communications channels on trust development. (2002), 135.

[4] Campos, P.F. 2004. *The obesity myth: Why America's obsession with weight is hazardous to your health*. Penguin.

[5] Cardiogram for Apple Watch and Android Wear: what's your heart telling you? *http://cardiogr.am/*. Accessed: 2016-05-26.

[6] Cheshire, C., Gerbasi, A. and Cook, K.S. 2010. Trust and Transitions in Modes of Exchange. *Social Psychology Quarterly*. (Jan. 2010).

[7] Cook, K.S., Yamagishi, T., Cheshire, C., Cooper, R., Matsuda, M. and Mashima, R. 2005. Trust Building via Risk Taking: A Cross-Societal Experiment. *Social Psychology Quarterly*. 68, 2 (Jun. 2005), 121–142.

[8] Cronin, B. and Sugimoto, C.R. 2015. *Scholarly metrics under the microscope: from citation analysis to academic auditing*. Information Today.

[9] Curmi, F., Ferrario, M.A., Southern, J. and Whittle, J. 2013. HeartLink: open broadcast of live biometric data to social networks. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), 1749–1758.

[10] Daft, R.L. and Lengel, R.H. 1984. Information richness: A new approach to managerial behavior and organizational design. *Research in Organizational Behavior*. (1984).

[11] Decaria, M.D., Proctor, S. and Malloy, T.E. 1974. The effect of false heart rate feedback on self-reports of anxiety and on actual heart rate. *Behaviour Research and Therapy*. 12, 3 (Sep. 1974), 251–253.

[12] Donath, J. 2007. Signals in social supernets. *Journal of Computer-Mediated Communication*. 13, 1 (2007), 231–251.

[13] Elsden, C., Mellor, S., Olivier, P., Wheldon, P., Kirk, D. and Comber, R. 2016. ResViz: Politics and Design Issues in Visualizing Academic Metrics. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), 5015–5027.

[14] Frey, J. 2016. Remote Heart Rate Sensing and Projection to Renew Traditional Board Games and Foster Social Interactions. *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (2016), 1865–1871.

[15] Howell, N., Devendorf, L., Tian, R.K., Galvez, T.V., Gong, N.-W., Poupyrev, I., Paulos, E. and Ryokai, K. 2016. Biosignals as Social Cues: Ambiguity and Emotional Interpretation in Social Displays of Skin Conductance. *Proceedings of the Designing Interactive Systems Conference* (2016).

[16] Janssen, J.H., Bailenson, J.N., IJsselsteijn, W.A. and Westerink, J.H.D.M. 2010. Intimate Heartbeats: Opportunities for Affective Communication Technology. *IEEE Transactions on Affective Computing*. 1, 2 (Jul. 2010), 72–80.

[17] Kastrenakes, J. 2014. Apple Watch uses four sensors to detect your pulse. *The Verge*.

[18] Lotan, G. and Croft, C. 2007. imPulse. (2007), 1983.

[19] McNell, J. 2015. Who Sexts Thumbprints? *Medium*.

[20] Merrill, N. and Cheshire, C. 2016. Habits of the Heart(rate): Social Interpretation of Biosignals in Two Interaction Contexts. (2016).

[21] Min, H.C. and Nam, T.-J. 2014. Biosignal sharing for affective connectedness. (2014), 2191–2196.

[22] Parkinson, B. 1985. Emotional effects of false automatic feedback. *Psychological Bulletin*. 98, 3 (1985), 471–494.

[23] Reeves, S., Martindale, S., Tennent, P., Benford, S., Marshall, J. and Walker, B. 2015. The challenges of using biodata in promotional filmmaking. *ACM Transactions on Computer-Human Interaction (TOCHI)*. 22, 3 (2015), 11.

[24] Scissors, L.E., Gill, A.J., Geraghty, K. and Gergle, D. 2009. In CMC we trust: the role of similarity. (2009), 527.

[25] Slovák, P., Janssen, J. and Fitzpatrick, G. 2012. Understanding heart rate sharing: towards unpacking physiosocial space. (2012), 859.

[26] Spire is the first wearable to track body, breath, and state of mind.: *http://www.spire.io*. Accessed: 2016-05-26.

[27] The best biometric and heart rate monitoring headphones: *http://www.wareable.com/headphones/best-sports-headphones*. Accessed: 2016-05-22.

[28] Valins, S. 1966. Cognitive effects of false heart-rate feedback. *Journal of Personality and Social Psychology*. 4, 4 (1966), 400–408.

[29] Vermeulen, J., MacDonald, L., Schöning, J., Beale, R. and Carpendale, S. 2016. Heartefacts: Augmenting Mobile Video Sharing Using Wrist-Worn Heart Rate Sensors. *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (2016), 712–723.

[30] Werner, J., Wettach, R. and Hornecker, E. 2008. United-pulse: feeling your partner's pulse. (2008), 535.

[31] Wilson, C. and Oeppen, J. 2003. On reification in demography. *Population, projections and politics*. (2003), 113–129.

[32] Wu, L., Waber, B., Aral, S., Brynjolfsson, E. and Pentland, A. 2008. Mining face-to-face interaction networks using sociometric badges: Predicting productivity in an it configuration task. (2008).

[33] Young, D., Hirschman, R. and Clark, M. 1982. Nonveridical heart rate feedback and emotional attribution. *Bulletin of the Psychonomic Society*. 20, 6 (Dec. 1982), 301–304.