

Shopping for Confirmation: How Disconfirming Feedback Shapes Social Networks

Paul Green, Jr.
Francesca Gino
Bradley Staats

Working Paper 18-028



Shopping for Confirmation: How Disconfirming Feedback Shapes Social Networks

Paul Green, Jr.
Harvard University

Francesca Gino
Harvard Business School

Bradley Staats
University of North Carolina

Working Paper 18-028

Copyright © 2017 by Paul Green, Jr., Francesca Gino, and Bradley Staats

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Running Head: SHOPPING FOR CONFIRMATION

**Shopping for Confirmation:
How Disconfirming Feedback Shapes Social Networks**

Paul Green, Jr.

Harvard University

Francesca Gino

Harvard University

Bradley Staats

University of North Carolina

September 15, 2017

Acknowledgements

We are grateful for the assistance and support of many individuals at our field site. In particular, we are grateful to Ron Caoua for his patience and support in gathering the rich field data that forms the backbone of this paper. We also thank Hayley Blunden, Joshua Margolis, Sharon Parker, Peter Marsden, and the WOM group and OB lab at Harvard Business School for their valuable comments on prior versions of this manuscript.

ABSTRACT

Many organizations employ interpersonal feedback processes as a structured means of informing and motivating employee improvement. Ample evidence suggests that these feedback processes are largely ineffective, and despite a wealth of prescriptive literature, these processes often fail to lead to employee motivation or improvement. We propose that these feedback processes are often ineffective because they represent threats to recipients' positive self-concept. Because the self-concept is socially sustained, recipients will flee these threats, or otherwise reshape their network to attenuate the negative psychological effects of the threat. Analyzing four years of peer feedback and social network data from an agribusiness company in the Western U.S., we find that employees, in the face of feedback that is more negative than their own self-assessment in a given domain (i.e., disconfirming feedback), reshape their network in ways designed to attenuate the threat brought about by the feedback, and that this behavior is detrimental to their performance. In a laboratory study, we replicate these findings conceptually, showing that disconfirming feedback has such effects on one's relationships and performance because it is perceived as threatening to one's self-concept.

Keywords: Developmental feedback, Self-concept, Positive illusions, Social network, Threat

INTRODUCTION

In any competitive organizational environment, enhancing employee performance is critical to sustained performance. Organizations seek ways to enable employee learning and development and to spark innovation and incite motivation, employing myriad organizational systems aimed simply at employee improvement. Organizations commonly use feedback processes as a way to provide employees with insight into their own work behavior and provide an impetus for improvement (Murphy and Cleveland, 1995), but though performance feedback processes—organizational processes that include a self-evaluation combined with peer-evaluations—are common across many organizations, and have been long, and widely, studied by organizational researchers, an uneasy consensus has arisen that interpersonal feedback processes have not lived up to their billing, with some scholars calling for their immediate abandonment, referring to performance feedback as a “failed experiment” (Adler et al., 2016).

In this paper, we address calls for critically executed research on the effects of developmental feedback on employees’ performance and relationships at work by exploring how employees experience such feedback. Though most prior research in this domain focuses on enabling effective improvement through feedback, we adopt a different focus, and investigate one previously unstudied function of interpersonal developmental feedback. Drawing on social psychology research on self-protection (Alicke and Sedikides, 2009; Sedikides and Alicke, 2012), we argue that employees look to peer feedback for confirmation of their self-concept. In addition to (or instead of) developmental insight, employees look to peer ratings as signals about relational others’ view of the self. In brief, peer feedback processes are stylized relational interactions, replete with clear signals to the recipient about how the feedback givers view the recipient. Employees experience feedback that is relatively lower than their self-views as a threat to their self-concept. We argue and show that employees will flee the threat brought on by

feedback, severing relationships with feedback givers when possible. Because a primary psychological purpose of relationships is to provide affirmation of the self-concept, employees respond to these threats by reshaping their social context in search of a suite of relationships that they expect will serve to better sustain their self-concept into the future. This behavior—the reshaping of the social network in search of more confirmatory relationships—has unfortunate negative side effects on subsequent performance. We test our theory in a field study, using four years of peer feedback and social network data collected from a vertically integrated agribusiness and food processing company, as well as in a laboratory study, providing causal evidence for our hypotheses and allowing us to examine the psychological mechanism explaining the relationship between disconfirming peer feedback and performance.

This paper makes several contributions to the literature. First, we contribute to existing feedback research by examining how feedback that does not confirm employees' positive self-concept influences their willingness to maintain social ties with feedback-givers. Specifically, we argue that though feedback interventions are motivated by a unitary purpose—enabling and facilitating improvement—these interventions are also interactions in which the recipient receives information about how they are viewed by the feedback giver. Although some recent research has explored the psychological effects of negative feedback, this work has generally examined the effects of such feedback on performance in the task domain without considering the broader behavioral effects of peer feedback (Ilgen and Davis, 2000; Greve, 2003; Jordan and Audia, 2012). Here, we suggest that, because an individual's self-concept is maintained and validated through the social environment, negative feedback will have social effects resulting from the individual drive to cultivate a socially validating environment.

Second, our research identifies a previously unexplored explanation for why feedback may not yield performance improvements. We argue and find that feedback that disconfirms

one's self-assessment is perceived as threatening to the recipient's self-concept, and that employees attempt to mitigate its psychological effects by altering their social network. Because relationships may be more readily exchanged in organizational settings than job tasks or activities, employees ultimately may cultivate a social network that does not provide adequate support for their suite of job responsibilities. Employees engage in an unfortunate exchange: trading away access to social resources that might improve their performance in the long-term in exchange for a relationally bolstered self-concept.

Finally, we provide the first (to our knowledge) empirical examination of the relationship between feedback and organizational social networks. Drawing on psychological research on positive self-concept maintenance, we argue that employees may react to disconfirming feedback by cutting ties with feedback providers, thus reshaping their social networks over time. This insight shines light on an important and unexplored domain in social network research, adding to the growing examination of network dynamics, exploring the predictors of network position (Burt, Jannotta, and Mahoney, 1998; Mehra, Kilduff, and Brass, 2001; Kadushin, 2002; Borgatti and Foster, 2003; Rivera, Soderstrom, and Uzzi, 2010; Kleinbaum, Stuart, and Tushman, 2013; Ma, 2015). Of potentially greater import, we argue for an expanded view of social capital in interpersonal networks, proposing that individuals look to their social environment as a source of self-concept validation, in addition to more traditional forms of social capital (such as social support, and exclusive access to unique opportunities and information). This expanded view of social capital compels a more nuanced view of the inherent value in social capital, suggesting that all social capital is not created equal. Specifically, in organizational contexts, when individuals arrive at advantageous positions through the pursuit of self-confirmation, they may not be able to effectively capture the benefits of that advantageous network position, and may even perform more poorly, despite their relatively advantageous position. We believe this

expanded view of social capital, and the specific motivation for accruing social capital, hold great promise for the examination of individual-level networks, particularly in organizational settings.

THEORY

The Unfulfilled Promise of Peer Feedback Interventions

Feedback is generally characterized as a process of informing and enabling personal improvement and development. The feedback literature arose, in large part, from the realization that employee self-assessments of performance tend to be unrealistically upwardly biased and not particularly useful for developmental purposes. Personal improvement is generally thought to require a constant evaluation of one's deficiencies and a focus on positive change (Carver and Scheier, 1982). Yet, as humans, we tend to see ourselves in a flattering light, emphasizing the positive aspects the self, and discounting the negative aspects (Kunda, 1987; Vaillant, 1995; Sherman and Cohen, 2006). Indeed, Meyer (1980) described a number of studies finding that employees routinely rate themselves highly relative to their colleagues and above average for their organization. In one exemplar study, the average self-appraisal was the 78th percentile, with only two out of 92 total participants assessing themselves as below the 50th percentile (Kay, Meyer, and French, 1965). As these results indicate, we humans maintain unrealistically positive self-views (Greenwald, 1980; Taylor and Brown, 1988).

Performance feedback interventions are rooted in the assumption that, though employees may not be accurate self-appraisers, they can gain relatively more accurate insight into their actual behavior and performance via *others'* relatively negative assessments of them (Campbell and Lee, 1988). Feedback from others is not subject to the self-deceptive tendencies so apparent in the self-evaluation process. It is more likely to be honest and thus can provide insights that enable recipients to overcome ego-driven self-enhancement. When coupled with a self-

assessment, feedback from others is thought to illuminate discrepancies between actual and desired performance levels (Campbell and Lee, 1988; Mohrman Jr et al., 1989), in turn increasing the motivation to improve (London and Smither, 1995). Most performance management systems now incorporate some form of peer feedback element (Pulakos et al., 2015), either direct (e.g. from peer to peer) or indirect (e.g. collected and summarized by the recipients' manager). But although peer feedback can hold a mirror up to employees, identifying areas in need of development, this exposure may not lead to meaningful efforts to improve, even when the feedback is valuable.

Various empirical studies have found that employees' actual responses to negative feedback are complex; feedback interventions designed to illuminate employees' blind spots very often do not yield the desired outcomes (see Adler et al., 2016 for a review). Though feedback from others is intended to motivate improvement, in reality, these processes are often demotivating, even to high-performing employees (Aguinis, Joo, and Gottfredson, 2011; Culbertson, Henning, and Payne, 2013). One meta-analysis found that one-third of feedback interventions actually resulted in *lower* post-feedback performance (Kluger and DeNisi, 1996).

These observations about the general inefficacies inherent in feedback interventions have led to extensive empirical research, all striving to understand and rectify these issues, and to understand the antecedents or conditions enabling improvement as a result of feedback (Posthuma and Campion, 2008; Adler et al., 2016). The historical work, together, points to four general diagnoses for ineffective feedback interventions: poorly designed feedback instruments and measures; poorly executed feedback; contextual features that restrain feedback effectiveness; and psychological processes that distort negative feedback.

Poorly Designed Instruments. Research exploring the design features of feedback interventions tend to focus on structural aspects of the feedback process that influence the

reliability of the feedback. This research emphasizes measurement issues as a detractor from feedback effectiveness, pointing to interrater disagreement (Viswesvaran, Ones, and Schmidt, 1996; Murphy, Cleveland, and Mohler, 2001) and measures of rater error (Murphy and DeShon, 2000) as evidence of lack of construct validity (Ones, Viswesvaran, and Schmidt, 2008; Adler et al., 2016). A separate vein of study exploring design features influencing feedback effectiveness points to the confounding effects of conflicting purposes. Though feedback processes are, generally, designed to inform development and improvement, the data collected in the feedback process often serves double-duty, informing decisions such as compensation and promotion (Murphy and Cleveland, 1995). These often-conflicting purposes tend to pervert the feedback given, leading to chronic and systematic overrating (Cleveland, Murphy, and Williams, 1989). This line of study has received extensive attention over more than four decades (see, for example Heneman, 1974; Murphy, 1982; Pulakos, Schmitt, and Ostroff, 1986; Harris and Schaubroeck, 1988; Viswesvaran, Ones, and Schmidt, 1996). But more recent reviews and analyses suggest that, in part, the problem is less structural in nature (that is, not a function of construct validity) but rather one of different individuals, by virtue of their varying roles, having different perspectives of the feedback recipient (Murphy, Cleveland, and Mohler, 2001). In essence, the value in multi-rater feedback comes by having multiple individuals, with different perspectives, provide the focal employee with feedback. This is not a flaw, but rather a feature.

Poorly Executed Feedback. Research exploring the execution of feedback tends to emphasize interpersonal or execution issues influencing the efficacy of feedback. This work suggests that feedback is often ineffective either because (a) it is delivered poorly, and the message is garbled or confusing; or (b) in the case of numeric ratings, is not accompanied by behavioral cues that the recipient can draw on to improve their behavior (Murphy and Cleveland, 1995; Cannon and Witherspoon, 2005; Posthuma and Campion, 2008). This research points to

various feedback giver and receiver trainings, designed to help feedback givers accompany ratings with clear, behavior-based cues as to the reason for the ratings, and receivers develop behavior-based plans for improvement, as paths to improving feedback effectiveness (Day and Sulsky, 1995; Cannon and Witherspoon, 2005; Posthuma and Campion, 2008). Though these prescriptions have, in some cases, improved feedback effectiveness, on average, these interventions do not appear to have systematically improved feedback effectiveness in organizations (Murphy and Cleveland, 1991, 1995), likely because this line of intervention was based on a faulty diagnosis. Adler et al (2016: 225), in a recent review, suggest that there is “little real evidence that raters lack the ability to rate accurately,” and argue that the likely root cause has more to do with feedback recipients’ motivation and goals.

Contextual Features. Research examining the effects of contextual factors points to organizational features such as organizational culture and climate, or other group level factors such as trust, as enabling or restraining feedback effectiveness (Edmondson, 1999; Aguinis and Pierce, 2008; Ferris et al., 2008). Other authors point to socioeconomic factors (e.g. state of the national economy or labor market) as influencing feedback effectiveness (Grey and Kipnis, 1976; Murphy and Cleveland, 1991). This work, collectively, tends to suggest that contextual factors tend to influence feedback givers’ willingness to provide accurate feedback, leading to incomplete or inaccurate feedback.

It does appear that certain aspects of organizational or team climate can dramatically reduce the likelihood of learning or growth. This effect is driven in large part by feedback givers’ desire to avoid responsibility for negative consequences for feedback recipients. Creating a climate of safety, and eliminating any job or status threat from the feedback process can certainly increase the likelihood that feedback givers are more forthright in providing developmental feedback (London and Smither, 1995). But this line of exploration does little to

address the recipient side of the feedback equation: development as the result of feedback is certainly dependent on accurate, honest and forthright feedback, but ample evidence shows that feedback is often still not effective—even when the feedback provided is well-delivered, accurate and forthright (Murphy and Cleveland, 1995).

Though these three general streams of research are broad and extensive, and do seem to offer some strategies for improving feedback effectiveness, in aggregate, there is little evidence that feedback interventions, even employing best practices suggested by literature, have systematically lead to organizational level benefits (see DeNisi and Smith, 2014; and Adler et al., 2016 for reviews). The fourth, relatively more recent, line of examination points to a particularly promising avenue for understanding the ineffectiveness of peer feedback.

Identified Discrepancies Aren't Motivating. A final stream of research suggests that feedback does not improve performance because recipients, in light of information exposing a discrepancy between their self view and the views of others, simply revise their self-assessments downwards, or otherwise deny or ignore the feedback. Traditional theories of feedback research proposes that negative feedback illuminates a discrepancy between the desired state and the actual state—a discrepancy that the feedback recipient will be motivated to close (Tornow, 1993; London and Smither, 1995). Upon receiving negative feedback, this body of work suggests, recipients sometimes simply respond by changing their view of the self as reflected by lower self-assessments in future periods (Johnson and Ferstl, 1999), or by revising their performance aspirations downwards such that the actual performance is no longer satisfactory (Korsgaard, 1996). This revision of the aspirational self is often accompanied by some sort of rationalization process, in which the recipient looks to the social environment for permissive exemplars—others who also exhibit this lower standard, thereby permitting the lowered self-standard (Korsgaard, 1996).

More recent work in this vein of examination, while offering little in the way of prescriptions for maximizing feedback effectiveness, does point to existing work in social psychology suggestive of various psychological processes that might limit feedback effectiveness (Alicke and Sedikides, 2009; Sedikides and Alicke, 2012). Recent theoretical work draws on this social-psychological research to suggest that feedback interventions might not be effective simply because recipients ignore, reshape or otherwise avoid the discrepancy-illuminating information (Jordan and Audia, 2012). We extend this conceptual logic and identify a potential reason why negative feedback often fails to yield performance improvement: negative feedback is experienced as a threat to the self-concept, leading recipients to reshape their social network.

Disconfirming Feedback and Threats to the Self-Concept

The historical examinations of feedback effectiveness tend to ignore the multiplicity of human motives. In attending to the general human motive to improve, feedback interventions seem blind to the fact that humans are also motivated to see themselves in a positive light (Sedikides and Strube, 1995) and to defend their sense of self-worth from psychological threats to their positive self-view (Alicke and Sedikides, 2009; Sedikides and Alicke, 2012).

Negative feedback from peers is more than just negative insight about the self; it can be experienced as a threat to the self-concept—a threat that might activate self-protective tendencies. Traditional peer feedback processes are, by design, meant to expose the recipient to their blind spots. They generally begin with a thoughtful and introspective self-evaluation which forms the basis of comparison against which the feedback recipient compares feedback ratings they receive from others (Mohrman Jr et al., 1989; Korsgaard, 1996). But in so doing, they do not merely illuminate discrepancies between self and other-views, they also present a psychologically traumatic challenge to the recipient. One experimental examination of peer

feedback in the field is suggestive of this possibility, finding that recipients who first completed a self-evaluation were less satisfied with the feedback they received than those who did not complete a self-evaluation (Roberson et al., 1993). The authors attribute this overall satisfaction to the self-appraisal process: those who spend time reflecting on their past performance are likely more confident in the accuracy of their self-assessment than those who are not.

In short, all individuals are motivated to maintain a positive self-concept. But the self-evaluative process inherent to most feedback processes involves giving explicit voice to the precise numerical representation of that self-concept. Any subsequent peer feedback that is lower than the recipient's self-evaluation is experienced as *disconfirmatory*, in that, because it is relatively lower than the recipient's self-evaluation, the feedback challenges the recipient's explicit view of the self (Jordan and Audia, 2012). Because this view of the self is core to the recipient's self-esteem, this disconfirming evidence is seen as threatening to the psyche. Feedback that is equal to or more favorable than one's self-assessment as *confirmatory*, as it is seen as confirming (and perhaps even enhancing) one's positive views of the self (Jordan and Audia, 2012).

To this point in this manuscript we have, consistent with the bulk of the feedback literature, used the term negative in referring to feedback that exposes a discrepancy between the self-evaluation and the evaluation of others. Here, we draw an important distinction between negative feedback and feedback that is experienced by the recipient as disconfirming. Negative and disconfirming are orthogonal characteristics of feedback; negative resides in the mind of the feedback giver, and is a function of that giver's body of colleagues to whom they provide feedback, whereas disconfirming is a function of the received score relative to the self-evaluation. Recipients of feedback scores infer negativity by comparing the received score to their self-evaluation score. Feedback that is lower than the recipient's self-evaluation is

disconfirming; it challenges the employee's explicit self-concept, and is experienced as psychologically traumatic (Kay, Meyer, and French, 1965; Ilgen and Davis, 2000), particularly when on dimensions along which there is a ubiquitous preference for more positive ratings.

Recipients will will perceive this feedback as threatening to the self-concept. We thus predict:

Hypothesis 1: People are likely to perceive disconfirming feedback as more threatening to their self-concepts than feedback that is not disconfirming

The Socially Sustained Self-Concept

When we are presented with information about ourself that challenges our positive self-view, extensive social-psychological research shows that we engage in self-protective psychological processes—ignoring, minimizing, or reconstructing the information to the extent possible (Alicke and Sedikides, 2009; Sedikides and Alicke, 2012). For example, we soften negative information about ourselves by drawing positive comparisons to individuals who are worse on a given dimension (Wills, 1981) or by derogating others (Fein and Spencer, 1997). Other self-protective strategies include externalizing the causal attributions we make about our failures, such as finding flaws in a test that we fail (Wyer and Frey, 1983) or crediting an evaluator's negative feedback to that person's racism or sexism (Crocker et al., 1991). These insights suggest that, when presented with disconfirming feedback from others, we may respond with various psychological processes designed to mute the harmful effects of the threat.

But these self-protective strategies are purely psychological mechanisms for dealing with disconfirming feedback. They fail to account for the fact that our self-concept is socially maintained—that it is not enough for our positive view of the self to reside in our own minds; we need to maintain a critical network of relationships that, collectively, provide a sort of fuel for the maintenance of one's sense of self-worth (Swann and Read, 1981; Leary et al., 1995; Leary and Kowalski, 1997). Selective attention or self-deceptive tendencies are short-term palliative responses to the threatened self-concept, but are inherently incomplete accounts for an

individual's response to disconfirming feedback: though it is important that a person feel good about themselves, they also must ensure they are firmly ensconced in a social context that provides confirmation of their self-concept. In domains where the feedback comes by way of our social environment, and in domains that are central to our self-concept, we are constantly looking to construct a relational environment that confirms our self-concept—an insight presaged by Swann in his theory of self-verification (Swann Jr, 1983).

Indeed, validation and sustenance of the self-concept is one critical psychological motivation for forging and maintaining relationships with others (Leary et al., 1995; Leary and Kowalski, 1997). Ironically, feedback systems designed to provide information about the recipient to the recipient, serve to provide insight to the recipient about the feedback givers' view of the self. Even well-designed feedback interventions send powerful signals as to the giver's view of the recipient. Humans in any social setting actively monitor the social environment for cues as to how others view them, and a peer feedback process provides ready, and explicit, insight into others' views of the self (Leary et al., 1995).

In an organizational setting, we suggest, when people's self-concept is threatened, they will endeavor to shape their context to minimize or eliminate the threat. The threat to the self-concept posed by disconfirming feedback is social in nature. It is not enough to simply derogate or otherwise diminish the threat; because confirmation from others is the fuel that sustains the self-concept, employees must look for opportunities to reshape their network of work relationships in ways that they believe will be less threatening, and provide a critical mass of confirmation. Swann suggested that, in pursuit of a social environment that will provide vital confirmation of the self-concept, we attempt to identify relationship partners who we believe are likely to affirm the self (Swann Jr, 1983). Psychological response mechanisms in response to self-concept threats cannot effectively sustain the self-concept over the long-term. Employees

must, in the face of the psychological consequences of disconfirming feedback, find a way to ensure that their social environment provides the sustenance their self-concept craves. One path by which this may occur is through the strategic reshaping of one's social network in pursuit of future confirmation. Employees may strategically eliminate and/or add professional relationships in an attempt to eliminate (or minimize the risk of) disconfirming feedback and to ensure more favorable future feedback. In short, the psychological threat to the self-concept, brought on by disconfirming feedback, serves as a psychological nudge—an incentive, of sorts—to reshape one's set of relationships. Feedback systems, designed to motivate performance, by serving up a threat to the self-concept, can also motivate a shuffling of social networks.

We expect the nature of the strategic reshaping of one's social environment to be contingent on the nature of the recipient's relationship with the feedback giver. In many organizations, employees have some discretion over with whom they work, but certain aspects of their role will obligate them to maintain relationships with specific individuals. Employees will shape their social environment differently in response to disconfirming feedback from the former (a discretionary relationship) than they will to disconfirming feedback from the latter (an obligatory relationship). When possible to escape, we expect that employees experiencing a threat to the self-concept will attempt to flee the source of the threat. When employees receive disconfirming feedback from a colleague with whom they are not obliged to maintain a relationship by virtue of task or job related interdependencies or strictures, they will be motivated to eliminate that relationship to avoid the risk of threatening disconfirming feedback from that person in the future. We predict:

Hypothesis 2: People are more likely to eliminate a discretionary relationship with a person providing disconfirming feedback than they are to eliminate a discretionary relationship with a person providing feedback that is not disconfirming.

Hypothesis 3: Perceived threat to one's self-concept mediates the relationship between disconfirming feedback and the elimination of a discretionary relationship.

Of course, many relationships within organizations are not discretionary. Further, many peer review processes are anonymous; it might be difficult to distinguish who precisely provided the disconfirming feedback, leaving the recipient without the option of severing the relationship. We consider these relationships as obligatory and expect that employees will respond differently when receiving feedback from obligatory colleagues than when receiving feedback from discretionary colleagues. While we expect dropping a discretionary relationship to be sufficient to eliminate a threat to the self-concept, if an employee cannot simply drop a relationship in response to disconfirming feedback, we propose that they will seek new relationships elsewhere within the organization, specifically with colleagues from other areas who will be less likely to be influenced by the extant opinions of current colleagues. More importantly, the greater the number of obligatory relationships (colleagues in the recipient's core group of relationships) provide disconfirming feedback, the further away we would expect the recipient to look for these new relationships. This prediction is grounded in the logic that, if an employee cannot eliminate a disconfirming connection due to its obligatory nature, they will attempt to attenuate the future psychological effects of that disconfirming evidence in by shopping for new relationships that will (presumably) provide more favorable future reviews of the focal employee.

This logic aligns closely with the idea of constraint conceptualized by Burt (2009) in his seminal work on structural holes. Networks (or subsets of networks) can be described, in part, by the degree to which people are connected to others who are also connected to the same others. This work argues that non-redundancy of relationships (connections to others who aren't also connected to each other) yields a form of social capital. For example, consider a network with five actors. There are a total of 10 possible connections in this network—the number of connections that would arise if each actor is connected to each other actor in this network. Burt proposed that, while some positive aspects associated with such a network do exist, there are also

individual downsides. If an actor is connected to all other actors in a network, all of which are also connected to each other, the focal actor is relatively constrained by being so tightly enmeshed in a cohesive group of fellow actors (Burt, 2009). Interestingly, the relationship development patterns suggested by the social network literature—homophily and the forging of relationships with those to whom one is already relatively close—are suggestive of networks becoming increasingly constrained over time.

Conversely, if an actor is connected to all other actors, who are otherwise disconnected, the focal actor is relatively unconstrained, and has an advantageous network position, given that they are a path through which other actors must traverse in order to reach other connections in the network, and giving them unique access to information, opportunities and resources that come via non-redundant connections. Burt developed an egocentric network index to measure the degree to which an individual is constrained, ranging from 0 (unconstrained, in which the focal actor is connected to all others in the network, who are otherwise not connected to each other), to 1 (constrained, in which all actors in the network are connected to all other actors).

When employees receive disconfirming feedback from obligatory relationships, employees will seek out new relationships with employees outside of their tight-knit circle of colleagues who are the source of the troubling disconfirming reviews. The greater the number of obligatory peers provide disconfirming feedback, the further away (and more relatively disconnected from the recipient's core group of connections) the new connection will be—leading to relatively lower constraint. Thus, we predict that:

Hypothesis 4: The greater the number of a person's obligatory reviews are disconfirming, the greater the negative change in future constraint.

Shopping for Confirmation and Performance

A healthy self-concept is critical to the maintenance of general psychological well-being and avoidance of depression and negative affect (Ruehlman, West, and Pasahow, 1985; Taylor

and Brown, 1988). The indirect organizational effects associated with perceived threats to an employee's self-concept are multitude—ranging from disengagement to turnover to decreased motivation (Elsbach, 2003; Davies, Spencer, and Steele, 2005; Nag, Corley, and Gioia, 2007). However, this leads us to the tension that prompted feedback interventions in the first place: we are willfully blind to our deficiencies, and improvement demands some sort of outside intervention to make visible the areas needing improvement. An individual who indulges the drive to reshape their network of work relationships as a means of avoiding a threat to their self-concept might consequently be able to maintain their motivation and sense of engagement, but by eliminating discretionary relationships, may be less likely to receive the support and advice that is key to one's career and development. Work performance depends, to some degree, on obtaining the information needed to solve challenging problems as well as advice from others, who can provide different perspectives on the task and appropriate support along the way (Mehra, Kilduff, and Brass, 2001; Cummings and Cross, 2003). The network of colleagues who provide developmental feedback is also a source of valuable resources—enabling the transfer of institutional knowledge, critically valuable advice and problem-solving assistance (Brass, 1981; Baldwin, Bedell, and Johnson, 1997; Sparrowe et al., 2001; Reagans and McEvily, 2003; Cross and Cummings, 2004; Gino, 2008). Dropping relationships in response to disconfirming feedback robs the recipient of potentially valuable sources of advice, assistance, and overall support. Thus, we hypothesize:

Hypothesis 5a: Eliminating discretionary relationships with individuals who provided disconfirming feedback is negatively associated with subsequent performance.

Even seeking distant relationships as a response to disconfirming feedback from obligatory relationships is likely also a reflection of the recipient's motivation to distance themselves from the disconfirming feedback giver. We argue with our Hypothesis 4 that increased incidence of received disconfirming reviews by obligatory relationships will yield a decrease in

the recipients' future constraint as the recipients seeks ever-more distant relationships that they hope have the prospect of future confirmation. This assertion exposes an important tension. Burt (2009) conceptualized constraint as an important indicator of social capital. Actors who are relatively less constrained hold a relatively unique spot within their network, spanning portions of the overall network that are otherwise relatively disparate and unconnected. Consequently, lower constrained actors are thought to have unique access to a more diverse set of informational and talent resources, as well as unique exposure to a broader variety of opportunities (Burt, 2009). Indeed, most research examining the affects of constraint (or other similar measures of social capital) suggests substantive benefits associated with decreased constraint (Burt, Jannotta, and Mahoney, 1998; Mehra, Kilduff, and Brass, 2001; Rivera, Soderstrom, and Uzzi, 2010). The social capital literature, though, operates under the assumption that individual actors strategically forge relationships with disparate, otherwise unconnected, others in pursuit of these informational and resource advantages (Marsden, 1987; Burt, 2005, 2009).

Our arguments suggest, though, that employees engage in strategic network formation in pursuit of a broader set of social benefits—that they are motivated to forge non-redundant relationships in pursuit of confirmation, a unique form of social capital. The organizational benefits of social capital are obvious: more diverse information sources, a broader set of resources and opportunities increases the likelihood of explorative innovation, leading to the sorts of innovation and entrepreneurial advantages central to Burt's (2009) theory. But at the individual level, social capital also includes the social sustenance of the self-concept; that is, constraint reductions emerge in response to disconfirming feedback not merely because employees are seeking unique informational or opportunity advantages, but because they are seeking social confirmation of the self. At the individual level, the basket of socially facilitated benefits must be expanded to include bolstering and support of the self-concept.

Paradoxically though, it is likely that a seeker who finds a non-redundant relationship elsewhere in the organization in response to a challenged self-concept, has found a relationship that, though in a different “neighborhood” (leading to lowered constraint), is actually similar in thinking and viewpoint to the focal actor. Confirmation shopping leading to lower constraint leads to the formation of an “echo chamber” of sorts. And though this behavior leads to an egocentric network that assumes an increase in social capital, the form of the social capital is likely to come without the traditionally expected benefits of lowered constraints: a like-minded and relatively similar connection, even if in a disparate section of the organization, is unlikely to provide the depth of informational and opportunity access advantages that accrue to those who seek non-redundant relationships in pursuit of those valuable resources. In short, the motivation leading to lowered constraint is a critical moderator of the distinct effects of decreased constraint in an organizational setting, and when examining individual-level networks.

We propose that individuals who engage in confirmation shopping leading to relatively lower constraint will perform at lower levels relative to individuals who refrain from this constraint-lowering confirmation shopping behavior, even in light of disconfirming feedback. Thus, we hypothesize:

Hypothesis 5b: Decreases in constraint in response to disconfirming feedback by obligatory relationships is negatively associated with subsequent performance.

At a more general level, relationships can change more easily, and more rapidly, than job tasks. When employees respond to disconfirming feedback by reshaping their network they have, in a way, altered the suite of advisors and valued sources of social assistance in completing aspects of their job. They eliminate sources of advice and help necessary for them to adequately perform the task requirements of their role. Even in seeking new, relatively distant, connections, they introduce a wider, more diverse, set of relationships that they must now manage, straining

their relational resources, and definitionally requiring that they divert valuable energy away from relationships that could prove vital and valuable to the performing of their job, toward relationships that serve the primary purpose of sustaining the employee's positive self-concept. On the whole, this image, of a feedback recipient now managing a set of relationships relatively less well-suited to supporting their particular set of activities, supports the intuitive notion that employees may trade valuable task support (which can improve performance) for psychological support when they reshape their social network.

Overview of the Present Research

Figure 1 depicts our theoretical model. To test our hypotheses, we use both field data (Study 1) and a laboratory study (Study 2). In Study 1, we use multiple years of data from a peer feedback process, combined with a unique longitudinal social network dataset from the same organization, to test Hypotheses 2, 4, 5a, and 5b. We find that employees, in the face of disconfirming peer feedback, reshape their social networks in ways designed to eliminate or attenuate the threat brought about by the feedback, and that this behavior is detrimental to employee performance. In Study 2, we provide a conceptual replication of these findings and show that this reshaping of social networks occurs because negative feedback is perceived as threatening to one's self-concept, thus providing evidence consistent with Hypotheses 1, 2 and 3.

 Insert Figure 1 about here

STUDY 1: FIELD STUDY

Sample and Data

To examine Hypotheses 2, 4, 5a, and 5b, we use archival data collected over a period of four years from a vertically integrated food manufacturing and agribusiness company located in the Western United States. The firm has been in business for nearly 30 years and encompasses a

wide-range of interrelated functions, including farming, harvesting, trucking, processing and distribution. Over the period covered by our sample, the organization employed between 260–290 full-time employees at any given time (although, with attrition, hiring, and other staffing fluctuations, there are a total of 347 unique employees represented in our data).

The company operates under a fluid structure; employees have some discretion over the scope, responsibilities, and deliverables of their role. As part of this process, employees enter into agreements with other employees to ensure the coordination of effort, transfer of information, and successful flow of work through the organization. Individuals identify specific tasks or projects in which they wish to be involved and identify relationships with other employees who depend on or are involved in these projects or tasks. Some of these relationships are obligatory (e.g., two employees share a job function or are in the same functional area), while some are discretionary, and reflect projects and particular initiatives that an employee chooses to undertake. These colleague relationships form a social network of sorts, which we leverage in the testing of our hypotheses. Employees also are responsible for providing feedback to all of their identified colleagues, as the organization relies heavily on peer regulation and proactive self-development. This peer feedback process is the sole structural method of providing employees with developmental insight within this organization, as it has no managerial review processes.

Peer-Review Data

Each December, the entire company engages in a structured feedback process. Employees complete a self-evaluation, and provide structured reviews of each of their identified colleagues from the past year. Each employee is blind to others' reviews of them until after the review period closes. The self and peer-evaluations are identical and are completed using a proprietary software that includes Likert-type scales across seven organizationally important distinct categories (such as “Improvement Orientation” and “Leadership & Initiative”). The

individual items in each category were crafted and refined by the organization over a number of years by employees working in employee development in concert with researchers. The individual constructs, and in many cases, the specific items, have been discussed in open forums with broad representation of the employee population as a means of ensuring that the collective instrument is representative of the general, and aggregate, aspiration of the average employee within the organization.¹ They include items such as “[employee] proposes and causes changes in the way things are done in order to improve the operation” in the “Improvement Orientation” category and “[employee] takes unpopular stands when necessary and advocates for what they believe is right” in the “Leadership & Initiative” category. All items are rated using a 7-point Likert scale ranging from “Very Poor” to “Excellent” (as well as “Not capable of judging”). Each category includes optional text response fields, and though comments are not required, employees are encouraged to provide both constructive and positive comments.

Each year, prior to the review period, employees are invited to a series of seminars held at each location, in which a trainer describes the overall process, and provides advice and encouragement on the effective use of the process. Specifically, employees are asked to provide actionable and descriptive written feedback to their peers, particularly if the feedback score is relatively low. Each employee is asked to set up a time to meet with each colleague providing a review after the review period is closed, to discuss the review, and to discuss strategies for improvement. At the end of the review period, each employee receives a report that includes a summary of their self-assessment on each item (completed prior to seeing the reviews provided by their colleagues), each received review (not anonymized), and an averaged summary of all of

¹ We note that the exact items changed over the course of the four years. The most notable difference is that, over the four-year period, the total number of items went from 34 to 20. The general categories remained stable, but the organization stripped away some items that seemed repetitive. The company aggregates the totals in the report; thus, we kept the same approach in analyzing the data. However, we note that there are seven distinct categories, and the items all load strongly into a single factor (even when the analysis is conducted separately for each year).

the employee's received reviews. The first page of the report provides the individual's overall average self-rating (calculated as the total average of all Likert-type items, collapsed into a single average), alongside the overall average received peer rating (on those same items) by each peer colleague. Because this overall aggregated average score on the Likert-rated items—for the self-evaluation and for each received peer evaluation—are the primary data presented on the first page of each employee's report, we use these same aggregated average scores in our analyses.

Insert Figure 2 and Table 1 about here

Our dataset includes all self-evaluations and peer evaluations conducted in the years 2012, 2013, and 2014. The average number of reviews received by year is shown in Figure 2, and correlations of all variables used in our analyses are shown in Table 1.

Employee Relationship Network Data

Each January, employees identify those colleagues with whom they will work in the coming year using a proprietary software program. Identification of these relationships is a participative form of crafting the employee's role within the organization—by identifying those with whom the employee expects to have a substantive working relationship. These relationships might take the form of shared responsibilities (two colleagues are co-committed to a specific activity), or interdependencies. They may also take the form of relationships around people desiring to work closely together—perhaps on a special project—or around a shared goal, though not necessarily interdependency or even meaningful working relationship. For example, one employee in a steam generation area formed a relationship with an employee responsible for regulatory issues—not around a functional interdependency, but because they both had an interest in reducing pollution—the former by virtue of his role operating large boilers; the latter by virtue of his role dealing with regulatory bodies. Employees have some discretion over the

structure of their role as well as the nature of their relationships within the organization. Some colleague relationships are discretionary (like the steam generation and regulatory relationship); others are obligatory, in that the physical process demands the two employees maintain a functional relationship (there are many interdependencies within the process, and committing to a specific activity may necessitate a relationship with a particular colleague). For example, if two employees are in the same work unit at a facility or location, it would be very difficult, given the structural interdependency that exists between the two, for them not to maintain a relationship.

Our dataset includes identified relationships for the entire company over four years: 2012, 2013, 2014, and 2015. The number of employees represented and the total number of identified relationships are shown by year in Figure 2. These datasets together provide review data collected over three consecutive years and employee relationship network data collected shortly thereafter (the subsequent January) for the following three years, as shown in Figure 2. This allows us to observe the ways in which the peer-evaluation process influences employees' social networks by observing changes in social networks in the period immediately following the review process.

Empirical Strategy

Disconfirming feedback and discretionary relationships. Hypothesis 2 predicted that employees would be more likely to drop relationships when, in the prior review period, they received disconfirming feedback from a colleague with whom they previously had a discretionary relationship than when they received feedback that was not disconfirming from a discretionary colleague. Because of our data's multi-level, nested structure—received reviews are nested within years, which are nested within employees—and because our dependent variable is dichotomous, we used Bernoulli distribution hierarchical linear models to test Hypothesis 2 (Raudenbush et al., 2016). We employ a three-level random intercept logistic

regression model, including random intercepts for each year (level two) and each feedback recipient (level three) to predict the likelihood of an employee dropping (eliminating a previously existing relationship) in year y , using the following model:

$$\begin{aligned} \text{DroppedNode}_{i,p,y} &= \beta_1 \text{Disconfirm}_{i,p,y-1} + \beta_2 \text{SameUnit}_{i,p,y} + \beta_3 \text{SameLocation}_{i,p,y} \\ &+ \beta_4 \text{TotalConfirmingReviews}_{i,y-1} + \beta_5 \text{ReviewerStdScore}_{i,p,y-1} \\ &+ \beta_6 \text{FeedbackGiverDisconfirm}_{p,i,y-1} + \beta_7 \text{Gender}_i + \beta_8 \text{SameGender}_{i,y} \\ &+ \beta_9 \text{SelfRating}_{i,y} + \beta_{10} \text{Tenure}_i + \varepsilon \end{aligned}$$

Where *DroppedNode* is a dummy variable, coded as 1 if, in year y , employee i eliminated a relationship that previously existed with peer p and is zero, otherwise; and where *Disconfirm* is a binary variable, coded as 0 if, in year y , peer p provided a rating of employee i that was greater than or equal to employee i 's self-rating, and coded as 1 if the received peer review was lower than employee i 's self-rating.

Controls

We include a number of control variables in our model. *SameUnit* is a dummy indicator that takes the value of 1 if both the focal employee i and the reviewing peer p are in the same work unit; *SameLocation* is a dummy variable that takes the value of 1 if both the focal employee i and the reviewing peer p work at the same location. These two variables together serve as a conservative proxy for obligatory relationships, as employees working in the same unit and at the same location are generally expected to maintain work relationships with each other.

TotalConfirmingReviews is a control for the total number of employee i 's received reviews in year $y-1$ that were not disconfirming. We would expect that the psychological threat to the self-concept would be more extreme for employees with only a few confirming colleague reviews than for employees with a great number of confirming reviews.

We also wish to control for the possibility that the severing of relationships is, at least in part, driven by the reviewer. Our analysis predicts the likelihood of a severed relationship. We presume a relationship to be severed by the feedback recipient, but in truth, relationships are truly dyadic; both parties must consent to the relationship. Consequently, a relationship could also be dropped by the reviewer; therefore, we must control for the prospective likelihood that the feedback giver is actually driving the dropping of relationships.

There are two relevant causes for a feedback giver to initiate the dissolution of a relationship. First, we would expect reviewers to drop relationships with those whom they rated lowest within their suite of provided reviews. We have argued that negative is in the eye of the feedback giver, whereas disconfirming is in the eye of the feedback recipient. As a reviewer, I am likely to want to discontinue my relationship with those in my network whom I perceive as the lowest performers—those to whom I give the relatively most negative rating. We would expect a reviewer's relative assessment of the recipient's overall performance to be reflected in the numeric rating of the recipient, and we would expect the feedback giver to be more likely to initiate the dissolution of a recipient whom they rate as lowest within their network of relationships. *ReviewerStdScore* controls for the review given to i by p , in year $y-1$, standardized by feedback giver p . Controlling for the review received standardized at the reviewer level effectively accounts for relative positivity/negativity of the review, by reviewer, and thus, the reviewer's likelihood of dropping the reviewee.

Second, a dropped relationship might occur because the feedback giver received disconfirming feedback from the recipient. *FeedbackGiverDisconfirm* is a dummy variable, taking the value of 1 if the focal employee i also provided disconfirming feedback to the feedback giver p in year $y-1$.

We also control for the recipient's gender using a dummy variable that takes the value of 1 if the focal employee i is female and is 0, otherwise. We also control for whether the feedback giver and recipient are the same gender using a dummy variable *SameGender* that takes the value 1 when feedback giver and recipient are the same gender and is 0, otherwise. We control for the recipient's self-evaluation score in year $y-1$ with *SelfRating*, and because we expect a recipient's tendency to respond to disconfirming feedback to vary as a function of their tenure with the organization, we control for tenure, a continuous variable measured in the number of years the recipient had been employed with the company at the time of the review.

Our Hypothesis 2 predicted that employees would be more likely to drop a discretionary relationship that provided disconfirmatory feedback in the current network period than they would be to drop a relationship that provided confirmatory feedback. This would lead us to expect that the odds ratio for *Disconfirm* will be greater than 1.

Disconfirming feedback and obligatory relationships. Hypothesis 4 predicted that the greater the number of an employee's reviews from obligatory relationships were disconfirming, the lower the employee's measured constraint in the subsequent network period. Given the time-series nature of our data, we analyzed the data using Hierarchical Linear Modeling (Hofmann, Griffin, and Gavin, 2000; Raudenbush and Bryk, 2002). We test Hypothesis 4 using a random intercepts model where the first-level model consisted of the change in constraint dependent variable regressed on the number of non-discretionary disconfirming reviews the employee received, as well as a number of controls. The second-level model consisted only of the randomly varying intercepts. We estimated the change in constraint for an employee from review year y to the subsequent network year, $y + 1$, using the following model:

$$\begin{aligned}
\text{ChangeConstraint}_{i,y,y+1} &= \beta_1 \text{CountNonDiscretionary}_{i,y} + \beta_2 \text{TotalDisconfirm}_{i,y} \\
&+ \beta_3 \text{TotalConfirmingReviews}_{i,y} + \beta_4 \text{ReviewerStdScore}_{i,y} + \beta_5 \text{SelfRating}_{i,y} \\
&+ \beta_6 \text{Gender}_i + \varepsilon
\end{aligned}$$

Where *ChangeConstraint* is the change in constraint from the review year, *y*, to the following network year, *y + 1*, for individual *i*, and where *CountNonDiscretionary* is the count of non-discretionary (obligatory) relationships providing a disconfirming review to employee *i* in year *y*.

We identified a relationship as non-discretionary if the fellow employee was in the same business unit at the same location as the focal employee. This is a conservative proxy for obligatory relationships in that, based on organizational rules, employees who work in the same area (or business unit, in this firm's lexicon) at the same location are required to work together.

Controls

We include a number of controls in our model. We control for *TotalDisconfirm*, the number of disconfirming reviews received by employee *i* in year *y*. Second, consistent with our specification for evaluating Hypothesis 2, we would expect that the psychological threat to the self-concept would be more extreme for employees with only a few confirming colleague reviews than for employees with a greater number of confirming reviews, and that those employees with a greater number of confirming reviews would be less threatened by a disconfirming review. That is, a single obligatory colleague providing a disconfirming review against the backdrop of 10 colleagues who have provided confirming reviews is likely to be less psychologically traumatic than that same single disconfirming obligatory colleague review against the backdrop of only a single confirming review. So we control for *TotalConfirmingReviews*, the total number of confirming received by employee *i* in year *y*. Also consistent with our specification for evaluating Hypothesis 2, we control for the average of *ReviewerStdScore* for employee *i* in year *y*, as well as a dummy variable of focal employee's

Gender and *SelfRating* for year y . We include year fixed effects to control for any systematic annual variance in shifting constraint.

We calculated each employee's egocentric constraint score, for each year, consistent with Burt (2009) using the formula below. Constraint c for focal employee i is the sum of the strength of each of i 's relationships with j , as a function of the proportion of i 's connections with all other contacts q are also relationships held by the direct contact j . In short, this function aggregates the proportion of each actor's relationships that are also mutually connected; a lower constraint score reflects a higher proportion of relationships that are not also mutually connected to each other (a relatively less constrained egocentric network).

$$c_i = \sum_j^{n_i} c_{ij} = \sum_j^n (p_{ij} + \sum_{q \neq j} p_{iq} p_{qj})^2$$

Our Hypothesis 4 predicted that, controlling for the total number of disconfirming reviews received by an employee, the greater the number of an employee's non-discretionary colleagues providing a disconfirming review, the further outside their focal network that employee would seek new relationships in future years, leading to reduced constraint. Thus, we expect that the coefficient on *CountNonDiscretionary* will be less than 0.

Results

Testing Hypothesis 2. As Hypothesis 2 predicts, we find a strong positive relationship between disconfirming feedback and the likelihood that the individual receiving the negative feedback drops the relationship in the subsequent year. The results are shown in Model 2 in Table 2. The statistically significant odds ratio of 1.367 ($p = 0.010$) on *Disconfirm* suggests that, controlling for whether the reviewer is in the same unit and at the same location as the recipient,

a reviewer who provides disconfirming feedback is more likely to be eliminated from the recipient's network in the following year than is a reviewer who provides confirming feedback.²

 Insert Table 2 about here

Testing Hypothesis 4. We predicted, with Hypothesis 4, that the greater the number of an employee's non-discretionary reviews that are disconfirming, the lower that employee's constraint in subsequent periods. Table 3 shows the results of our analysis of this hypothesis. Model 1 shows all controls regressed on our dependent variable of *ChangeConstraint*, excluding the independent variable. The results shown in Model 2 in Table 3 provide support for our Hypothesis 4. As predicted, the coefficient on *CountNonDiscretionary*, -0.0169, is negative and statistically significant ($p = 0.007$), suggesting that controlling for the total number of disconfirming reviews, each additional non-discretionary disconfirming review is predicted to reduce the employee's constraint score by 0.0169 in the subsequent year. When employees receive disconfirming feedback from peers within their work unit—relationships that they can't eliminate by virtue of structural aspects of the work—they may respond by seeking relationships with colleagues from outside their core network. To put this in perspective, the average constraint score in 2012 was .2866, suggesting that the average employee's constraint would drop by about 5.9% in response to a single disconfirming review by an obligatory colleague.

 Insert Table 3 about here

Testing Hypotheses 5a and 5b. We predicted that engaging in these behaviors—either dropping discretionary relationships with those who provided disconfirming feedback

² We note that, though we treat disconfirmation as a binary indicator, our results are robust, in direction and statistical significance, when using difference scores.

(Hypothesis 5a) or seeking new relationships with those who are otherwise relatively disconnected from one's current network, leading to decreased constraint (Hypothesis 5b)—would lead to lower performance. The organization employs a performance-based process for awarding bonuses annually, using individual performance on job-specific quantitative measures of performance as a guide for providing bonuses. Importantly, the bonus process is not related to the peer-evaluation process; bonuses are not in any way structurally related to peer evaluations, but are based solely on quantitative performance on key aspects of the employee's job (e.g. productivity metric targets), as well as quantifiable performance on any special projects (including innovation projects) the employee was involved in during the prior year. Further, we find that received bonus is only weakly correlated with average received review within the given year. So the best proxy for objective individual performance is the employee's received bonus. The organization provided data for bonuses awarded at the end of 2015. Each record in the dataset is comprised of the employee ID and a percentage bonus, calculated as the percentage of salary awarded as a performance bonus.

Because the organization only provided one year of performance data (bonuses at the end of 2015), our performance analyses are restricted to the period from the end of 2014 through the end of 2015. Our general performance hypothesis is that employees who, upon receiving disconfirming reviews, engage in these counterproductive network reshaping behaviors as they begin their next year's work, will perform more poorly over the following year and thus receive lower year-end bonuses. Our analyses use feedback data from the end of 2014, network data from 2015, and bonus allocations from the end of 2015. Figure 3 provides a timeline depicting the point at which each portion of data used in these analyses were generated.

Insert Figure 3 about here

Our first performance hypothesis (H5a) relates to the dropping of discretionary disconfirming relationships. We estimate the effect of a dropped disconfirming relationship on subsequent year bonus percentage using OLS to estimate the following model:

$$\begin{aligned} BonusLog_{i,2015} = & \beta_1 DroppedDiscrev_{i,2014} + \beta_2 TotalDisconfirm_{i,2014} \\ & + \beta_3 AvgRating_{i,2014} + \beta_4 AvgRating_{i,2015} + \beta_5 Gender_i + \beta_6 BusUnitID + \varepsilon \end{aligned}$$

Where $BonusLog_{i,2015}$ is the log-converted bonus amount from the end of 2015 (log-converted because of a severe positive skew in the bonus percentages). There were no negative bonus amounts in the raw data, and we added 1.00 to all raw bonus percentages prior to log-transformation so that zero bonus values were not undefined. Our independent variable of interest is $DroppedDiscrev_{i,2014}$, calculated as the total number of discretionary relationships that were dropped by employee i , following a disconfirming 2014 review. We include four controls: $TotalDisconfirm$ is the total number of disconfirming reviews received by employee i in the 2014 review period. And, although the ratings generated by the review process have no structural formulaic relationship with bonus received, there may be a relationship between the average review received at the end of 2014 or 2015 and the 2015 bonus received. Consequently, we also control for $AvgRating$, the average peer rating received by employee i at the end of 2015, and at the end of 2014, as well as the recipient's gender, to control for any gender effects associated with compensation. Finally, our specification includes business-unit fixed effects ($BusUnitID$) as units have varying norms around bonus percentages, and average bonus varies as a function of business unit. Standard errors are clustered at the business-unit level.

Hypothesis 5a suggests that the greater the extent to which individuals engage in the practice of dropping discretionary relationships that provide disconfirming reviews, the lower their performance will be in the subsequent year. We would expect, then, that the coefficient on $DroppedDiscrev$ would be negative. Model 2 in Table 4 shows the results of our analysis of H5a.

Although the estimated coefficient on *DroppedDiscrevi,2014* of -0.0021 is, in fact, negative (suggesting that each negative review leading to a dropped relationship has a negative impact on the subsequent year's log-bonus percentage), the calculated p-value ($p = 0.771$) is not statistically significant. Our data do not support Hypothesis 5a.

Insert Table 4 about here

Our second performance hypothesis (H5b) relates to the change in constraint we observe when employees receive disconfirming reviews from obligatory colleague relationships. Our analysis of Hypothesis 4 supports our assertion that individuals will, in response to disconfirming reviews from obligatory peers, seek connections with others who are relatively disconnected from their current circle of relationships, leading to a decrease in calculated constraint in the subsequent year. Our analysis of the performance effects associated with this behavior exploits the variance we observe in this behavior. Specifically, some employees seek new, relatively unconnected relationships when receiving disconfirming feedback from their obligatory relationships. Others demonstrate restraint, refraining from seeking these new, relatively unconnected relationships in spite of the disconfirming feedback from their obligatory relationships. We can create a “confirmation shopping” index by calculating the residuals from our H4 regression estimate (see Appendix A for a brief description of this approach). These residuals allow us to denote particularly extreme changes in constraint, as a function of the disconfirming reviews received by the employee, and is by construction, orthogonal to the various controls used to predict the change in constraint (Healy and Serafeim, 2015). In short, this approach provides an effective means of detecting extreme indulgence (or restraint) in confirmation shopping behavior as a response to disconfirming reviews. Recall that our estimate of H4 predicted the change in constraint of an employee in year y based on the number of

disconfirming reviews received by obligatory colleagues in year $y-1$. Further, the predicted change in constraint for a shopper is negative (that is, we expect those who over-indulge in this shopping for confirmation behavior to have lower constraint in the subsequent year). Thus, we would expect shoppers (those who over-indulge in this shopping behavior) to have relatively negative residuals and improvers to have relatively positive residuals.

To exhibit these relationships, consider the following hypothetical conceptual example. Employee A received three disconfirming reviews from obligatory colleagues in 2014. Suppose our regression model from H4 predicted a change in constraint, given three disconfirming reviews, of $-.20$ (it does not; we simply use these numbers for ease of illustration). If, in reality, Employee A's observed change in constraint was 0, our residual would be:

$$\text{Residual}_A = \text{Observed} - \text{Calculated} = 0.00 - (-.20) = .20$$

Employee B also received three disconfirming reviews from obligatory colleagues in 2014. Again, our regression model from H4 would predict a change in constraint of $-.20$. This employee over-indulged in confirmation shopping, and in the following year, the employee's constraint changed by $-.30$. In this case, our residual for H4 would be:

$$\text{Residual}_B = \text{Observed} - \text{Calculated} = -.30 - (-.20) = -.10$$

In short, the residuals from H4 provide a convenient distribution of greater (less) than expected changes in constraint, as a function of disconfirming reviews. Using this distribution of residuals, and on the logic that individuals engage in this behavior to varying degrees, we can estimate the degree to which engaging in, or restraining from, confirmation shopping influences an individual's performance over the course of the succeeding year.

We estimate the effect of an employee's calculated residual from our estimate of H4 on the employee's following year performance, using OLS to estimate the following model:

$$\begin{aligned}
 \text{BonusLog}_{i,2015} &= \beta_1 H4Residual_{i,2014} + \beta_2 AvgRating_{i,2014} + \beta_3 AvgRating_{i,2015} \\
 &+ \beta_4 Gender_i + \beta_5 BusUnitID + \varepsilon
 \end{aligned}$$

Where *BonusLog* is the log-transformed bonus percentage earned by employee *i* at the end of the year 2015, and where *H4residual* is the calculated residual value after estimating our primary model from our estimation of H4 above, for employee *i* at the end of 2014. Because we employ the residuals from the regression predicting change in constraint from 2013 to 2014, we only included employees who were employed at the end of 2013 and participated in both the 2013 and 2014 review processes (thus having a constraint score for both 2013 and 2014 and, in turn, a residual from the end of 2014), and who were still employed at the end of 2015. Further, because we employed residuals from H4, we need not include controls for the number of disconfirming reviews received as our residuals effectively account for this variable. We control for *AvgRating2014*, *AvgRating2015*, and *Gender* (all consistent with our analysis of H5a). Our specification includes business-unit fixed effects (*BusUnitID*), with standard errors clustered at the business-unit level.

Our primary specification, shown in Model 4 in Table 4, supports our Hypothesis 5B. The calculated coefficient of 0.0320 is directionally as we would expect: a more positive residual, expected for an “improver,” yields a higher log-bonus; a more negative residual, expected for a “shopper,” yields a lower log-bonus, and is statistically significant ($p = 0.017$). Importantly, we find that overall constraint at the end of 2014 is positively associated with log-bonus at the end of 2015; specifically, we find a significant coefficient of -0.13516 ($p=0.000$) of 2014 Constraint on 2015 log-bonus, indicating that individuals who have a lower raw constraint score receive higher bonuses, controlling for business unit, average review score received, and gender, than those with higher raw constraint scores. Though not direct evidence, this suggests that the negative effects associated with lowered constraint are likely isolated to constraint that is

lowered as the result of the pursuit of confirmation of the self-concept. That is, constraint as an overarching construct behaves, in our field site, the way we would expect: lower constraint is associated with higher performance. It is only when individuals indulge in confirmation shopping, leading to reduced constraint scores, that we observe the adverse bonus effects.

Study 1 Discussion

The results of Study 1 provide support for Hypothesis 2 and suggest that disconfirming reviews lead employees to drop relationships with the reviewing colleague if the relationship is discretionary. The results also support Hypothesis 4: employees reform their network, likely in search of a more hospitable relational climate, by seeking new colleagues relatively more disconnected from their current circle of colleagues for each obligatory relationship providing a disconfirming review. Finally, we found partial support for our hypothesis that confirmation shopping leads to decreases in performance in the succeeding year. Our Hypothesis 5a was directionally as expected but not statistically significant. Our second performance hypothesis, H5b, was strongly supported, and suggests that those who indulge in seeking confirmation by finding relationships with those who are relatively unconnected from the focal employee's existing social network do in fact perform poorly in the subsequent year, relative to those who refrain from this confirmation-seeking behavior.

These findings provide some evidence that the network-shaping behavior that employees display in response to disconfirming feedback, while potentially salvaging their self-concept, is detrimental to long-term performance. Our performance measure is observed a full year after employees received disconfirming feedback and altered their network. A year is time enough for negative psychological effects of feedback to have been forgotten and for the recipient to have moderated their behavior and improved their performance. That we find such a performance effect even a year after the disconfirming feedback, suggests a more lasting mechanism

contributing to this decreased performance. In Study 2, we provide a conceptual replication of our field study, and validate the psychological mechanism leading to these behaviors.

STUDY 2: DISCONFIRMING FEEDBACK IS THREATENING

In Study 2, we examined the psychological mechanism explaining why disconfirming feedback leads to the reshaping of one's social network by focusing on the role of perceived threat to one's own self-concept (Hypotheses 1, 2, and 3). Two-hundred and three undergraduate and graduate students from a private university in the Northeastern United States (47.3 percent male; mean age = 21.67, s.d. = 1.34) participated in this study in exchange for \$10 and the opportunity to earn an additional \$5 bonus. Study 2 employed two between-subjects conditions: disconfirming vs. confirming feedback.

Procedure

The study was conducted online. The instructions informed participants that the researchers were interested in studying how people work with others and perform on creativity and problem-solving tasks. In the recruitment email, potential participants were told that the study was for people who value creativity and view themselves as creative.³ We used this wording to assure that only people who viewed creativity as an important aspect of their sense of self would complete the study, and feel threatened by the disconfirming feedback they'd receive in one of our conditions.

After answering questions about their gender and age, participants were told that, during the study, they would be paired with another participant and asked to work on a series of short tasks. After a short wait, during which the pairing was seemingly occurring, participants learned

³ Though we wanted to assure that only people who view creativity as an important part of who they are took the study so that the disconfirming feedback would be perceived as threatening, this does not seem to be an important feature of the study. In fact, we conducted the same study without this prompt in the recruitment and obtained the same results, in nature and significance.

that they would complete a creativity task with their partner. As the instructions explained, they would be randomly assigned to the role of either writer or evaluator in this task: “The writer will be asked to write a creative short story that is at least 200 words. The evaluator will be asked to evaluate it for creativity. You and the participant you have been paired with will now be randomly assigned to one of these two roles.” In reality, the other participant was a computerized script, and all study participants were assigned to the role of writers. To increase realism, the assignment to a partner and a task took some time to occur.

After being told that they had been assigned to the role of writer, participants were asked to indicate the extent to which they felt anxious, stressed out, and nervous ($\alpha = .96$) on a 7-point scale (from 1=not at all, to 7=very much so). We included this measure to assess their level of anxiety about the task they would soon be performing.

Participants next were given five minutes to write a story. Participants were told that the evaluator would be asked to evaluate the story for creativity. After writing their story, participants were asked to assess their own level of creativity based on what they had written (from 1=not creative at all to 10=very creative).

Next, participants were randomly assigned to one of two conditions: disconfirming vs. confirming feedback. In the confirming-feedback condition, they received a feedback score that was the same as their own self-assessment; in the disconfirming-feedback condition, they received a feedback score two points lower on creativity than their own self-assessment.⁴

Participants then answered a series of questions assessing how much they found the feedback to be threatening. We measured perceived threat using four items ($\alpha = .85$). Participants

⁴ We note that we conducted another study on a separate sample of students using the same task in which the confirming feedback was a score that was two points higher than participants' self-assessment. The nature and significance of the results did not change: receiving feedback that is either the same or above one's own self-assessment produced the same pattern of results as compared to the condition in which the feedback was disconfirming.

indicated the extent to which they found the evaluator's assessment to be: (1) threatening, (2) challenging, (3) disappointing, and (4) threatening to their sense of competence, each on a 7-point scale (from 1=not at all, to 7=very much).

Next, participants completed the Positive and Negative Affectivity Schedule (PANAS; Watson, Clark, Tellegen, 1988), a measure commonly used in the literature to assess emotions people are experiencing. For the PANAS, participants responded to both positive affect items (e.g., interested, excited, enthusiastic; $\alpha = .90$) and negative ones (e.g., distressed, upset, guilty; $\alpha = .94$), indicating the extent to which they felt each emotion in that moment on a 5-point scale (from 1 = very slightly or not at all to 5 = extremely). We also captured participants' level of anxiety after receiving the feedback by averaging three of the negative affect items, namely nervous, jittery and anxious ($\alpha = .87$) so that we could test whether anxiety served as an alternative explanation for the effect of disconfirming feedback on relationships.

Participants then moved to the next task: a trivia quiz. The instructions read:

In this task, you and another participant will be asked to answer a series of 10 questions under time pressure. If you both answer all questions correctly you will receive a bonus.

For this task, you can decide to continue working with the same participant you have been paired with or choose to be paired with a different one.

Prior to completing the trivia task, participants chose whether to complete the task with the same participant they had been assigned to at the beginning of the study or with another participant. This choice served as our main dependent variable. The instructions informed them that they and their partner would each respond to a set of 10 trivia questions and that if they both answered all the questions correctly, they would earn a \$5 bonus. After completing the trivia task, participants were debriefed. Each participant received \$15 for participating.

Results

Table 5 reports means and standard deviations, by condition, for all assessed variables.

Insert Table 5 about here

Perceived threat. Consistent with Hypothesis 1, participants who received disconfirming feedback perceived the feedback to be more threatening (mean = 2.67, s.d. = 1.38) than did those who received confirming feedback (mean = 1.54, s.d. = 0.88), $t(203) = 6.97, p < .001, d = .98$.

Choice of partner for the trivia task. A higher percentage of participants who received disconfirming feedback (29.7%, 30 out of 101) chose to be paired with a new partner for the trivia task as compared to those who received confirming feedback (8.8%, 9 out of 102), $\chi^2(1, N=203) = 14.25, p < .001, \text{Cramer's } V = .27$. These results are consistent with Hypothesis 2.

Anxiety. One may expect participants who received disconfirming feedback to feel higher levels of anxiety after receiving the feedback than those who received confirming feedback. This was not the case, $t(203) = 0.19, p = .85$. Similarly, the anxiety participants reported experiencing after the feedback as compared to the anxiety they felt prior to the creative writing task did not vary between conditions, $t(203) = 0.45, p = .66$.

Positive and Negative Affect. We also did not find differences in negative affect between conditions, $t(203) = 1.06, p = .29$. However, positive affect was higher for participants who received confirming feedback (mean = 3.04, s.d. = 0.84) than for those who received disconfirming feedback (mean = 2.67, s.d. = 0.84), $t(201) = 3.21, p = .002, d = .44$.

Performance on the trivia task. Given that the choice of partner did not affect work on the trivia task, we did not expect to find differences on performance. In fact, participants' score on the trivia task did not differ depending on whether they received disconfirming (mean = 4.39, s.d. = 1.95) or confirming feedback (mean = 4.64, s.d. = 2.17), $t(201) = -0.87, p = .39$.

Mediation analyses. Our third hypothesis predicted that perceived threat would explain why a disconfirming review of the focal person leads to the elimination of a discretionary

relationship. We tested whether perceived threat mediated the relationship between our feedback conditions and the choice of partner for the trivia task, using the bootstrapping approach outlined by Preacher and Hayes (2004). Based on bootstrapping (with 10,000 iterations), we estimated the direct and indirect effects of the feedback condition via perceived threat on our dependent variable: the choice to be paired with a new partner. Our manipulation had a significant effect on perceived threat (as shown by the analyses above), which, in turn, significantly affected the choice to be paired with a new partner ($B = .67$, $S.E. = .15$, $p < .001$). In contrast, the effect of our manipulation was reduced and became statistically not significant (from $B = 1.47$, $S.E. = .41$, $p < .001$ to $B = 0.74$, $S.E. = .46$, $p = .11$) when perceived threat was included in the equation. The 95 percent bias-corrected confidence interval for the size of the indirect effect excluded zero (.374, 1.281), suggesting that perceived threat mediated the link between the feedback condition and greater likelihood of choosing a different partner. We note that these results did not change when controlling for positive and negative affect, nor when controlling for the change in the level of anxiety participants experienced after receiving feedback.

Discussion

Together, these results provide support for Hypotheses 1, 2, and 3. We find that disconfirming feedback makes people less likely to interact with those who gave the feedback because they perceive it to be a threat to their self-concept.

GENERAL DISCUSSION

To improve and remain competitive, and in light of our self-deceptive tendencies, organizations employ developmental feedback processes designed to inspire greater individual effort toward personal growth and improvement. While in some cases, employees may embrace this feedback and work to improve, in a longitudinal field study and a lab study, we found that disconfirming peer feedback is experienced as a threat to the recipient's self-concept that leads to

a reshaping of their network. In Study 1, we found that employees are more likely to abandon relationships with colleagues who previously provided disconfirming feedback. When the disconfirming feedback comes from those with whom the employee must maintain a relationship, the employee is more likely to make connections in the subsequent year with other employees who are not densely connected to their current network of employees. Additionally, we found that those who engaged in such confirmation shopping, which leads to reduced constraint scores, experienced a significant drop in performance in the following year, suggesting that there is a performance cost associated with confirmation shopping: the strategic reshaping of an employee's social network in search of confirmation of the self-concept.

In Study 2, we conceptually replicated the findings in the lab, showing that individuals are more likely to drop relationships with those who provide disconfirming reviews than with those who provide confirming reviews and that this behavioral tendency is mediated by perceived threat to the self-concept.

Our findings provide further evidence that peer feedback processes are perhaps naively envisioned and that the logical consequence of developmental feedback is not necessarily individual development. We integrate concepts from identity threat, peer feedback, and social network research to demonstrate one harmful effect of developmental feedback, as well as the influence of socially activated psychological threat on network formation over time.

Theoretical Contributions and Practical Implications

This study highlights a critical tension between ensuring employees' well-being and psychological health, and providing them with critical developmental insight. Individuals do not necessarily rationally encode disconfirming feedback as a developmental insight; they can also experience it as a socially activated threat to their self-concept. We contribute to the feedback literature by demonstrating that disconfirming feedback yields a visceral psychological effect.

The more novel contribution, though, is that people's response to disconfirming feedback is not merely psychological in nature but also behavioral. Past research has argued and shown that people are, at times, feedback seekers (Ashford, 1986; Ashford, Blatt, and Walle, 2003). We suggest, too, that they are confirmation shoppers—and that they look to feedback systems as a domain in which they expect confirmation of the self. Importantly, the effects of this behavioral tendency—namely, the reshaping of one's social environment—are not confined to the domain of the feedback. That is, while derogation and discounting of negative feedback is targeted toward the source of the disconfirming feedback (Wyer and Frey, 1983; e.g. Crocker et al., 1991; Fein and Spencer, 1997), our findings suggest broader potentially harmful effects; responses to disconfirming feedback ripple out through the organization in the form of shifting networks.

We do not mean to suggest that all individuals universally seek to surround themselves exclusively with confirming relational others. Ample research across various domains point to various dimensions on which individuals may vary in terms of their need for relational self-confirmation. For example, an employee's implicit mindset (e.g. growth vs. fixed) might influence their need for confirmation (Dweck, 1986). Regulatory focus might influence the degree to which individuals seek or desire developmental insight, as promotion focused individuals are more concerned with the pursuit of the ideal self and are likely more able to brave the psychological pain associated with short-term disconfirmation, in service of long-term improvement than prevention focused individuals (Brockner and Higgins, 2001). The difference in the degree to which individuals feel threats to the self-concept in response to disconfirming feedback is likely a matter of degrees, though. That is, though there are almost certainly differences between individuals that lead to variance in sensitivity to disconfirming feedback, there is also a general need to maintain a core group of relationships that, together, provide

sustenance to the self-concept. We hope that our work will inspire further exploration of the individual differences that moderate the effects we observe in our data.

This research is not necessarily intended as a call for altering attributes of peer-feedback processes; much has been done on that front. Nor do we offer this study as evidence that even well-designed feedback systems can yield detrimental outcomes. On the contrary, though our partner organization's feedback process in some ways embraced best-practices, in other ways, it did not. Our intention is not necessarily to offer a mechanism by which feedback systems can be improved. Rather, we believe that feedback systems, and the behavior we observe in our data in response to feedback, provide a convenient means of demonstrating the importance of the relationally sustained self-concept. We view this research, first, as a call to recognize the relational implications of feedback interventions. Feedback is a relational cue, and in the absence of other sources of needed confirmation, individuals will look to feedback for this important socially-facilitated confirmation. Our discipline increasingly views the social fabric of the organization as a critical facet of the employee experience and of organizational efficacy (Dutton and Heaphy, 2003; Dutton and Ragins, 2007). Our work suggests that the performance systems we install within organizations can effectively rend that social fabric, causing deleterious shifts in organizational relationships that suppress many of the positive effects suggested by this broader body of research by unwittingly providing signals of the feedback giver's sentiments about the recipient—signals that can serve to challenge the recipient's sense of self. This research provides a first step toward a response to recent calls for feedback interventions that don't activate the "psychological immune system" (Neville and Roulin, 2016). This research speaks to the need to accompany developmental feedback processes with other means of bolstering the self-concept.

More generally, we view this work as suggestive of viewing employee improvement and development more holistically. Approaching feedback performance management interventions solely as an informational exercise is, we believe, a narrow, and likely ineffective, view. Practically speaking, our work suggests that the employee improvement endeavors may be conditional on the employee's healthy self-concept—a condition that seems at odds with a process designed, specifically, to undermine the self-concept. But relationships, viewed broadly, seem to hold the key to sustaining the self-concept, and we hope this work inspires further research exploring the ways in which relational features can increase employee receptiveness to disconfirming feedback.

Our work also contributes to social network research, both generally and as applied to organizations, much of which has been structuralist (Mayhew, 1980; Wellman, 1997) in nature (Borgatti and Foster, 2003). The structuralist paradigm tends to explore the effects of network structures on other outcome variables of interest and eschews the idea that individual differences can predict network characteristics (Mayhew, 1980). This study provides insight into an important way through which a key network characteristic emerges; constraint shifts, in part, in response to socially activated threats to the self-concept brought on by developmental feedback processes. Our data suggest that the search for a more hospitable set of professional relationships in the face of negative peer feedback yields a network characterized by higher levels of social capital but lower subsequent performance.

Relatedly, this research proposes an expanded view of social capital at the interpersonal level. Traditional conceptions of social capital view unique access to information, social support and opportunities as the valued currency flowing across relationships. We argue that individuals also value confirmation of the self-concept, and will strategically reshape their networks in search of distinct access to that relationally conveyed confirmation. Importantly, though advice,

information and unique opportunities generally lead to positive performance outcomes, a bolstered self-concept does not necessarily link to positive outcomes—and might, paradoxically, lead to negative outcomes, if the unique sources of self-confirmation are not the appropriate sources for advice, information and opportunities within the context of an individual's work role. This research suggests that individuals can seek, and occupy, advantageous network positions—but that the unique resources offered by those positions may not always lead to increased.

It is not lost on us that our findings suggesting a negative relationship between decreased constraint and subsequent year performance seem to challenge a widely-held tradition in the social network literature demonstrating that social capital yields benefits for the network occupant (Burt, Jannotta, and Mahoney, 1998; Mehra, Kilduff, and Brass, 2001; Sparrowe et al., 2001; Rivera, Soderstrom, and Uzzi, 2010). We do not have the data to determine whether, in the longer-term, the decreases in constraint observed in our data lead to performance increases, a possibility that we view as highly plausible. But neither do we mean for our study to challenge a long-standing tradition of research demonstrating real, and compelling, advantages to network position. Our research is a unique, panel view into networks, offering us the opportunity to attribute shifts in social capital, at least in part, to disconfirming feedback. We believe it possible, perhaps even likely, that in the longer-term, the decreases in constraint we observe will lead to the benefits so widely observed in much of the social network literature. But our data do suggest that, at a minimum, perhaps the benefits of social capital take time—they must accumulate, and are not readily accessible by the network occupant. In this way, we view our work as suggesting that, at the very minimum, in the short-term, social capital may have a cost.

Additionally, we see this research as adding to the growing research examining the ways in which organizational structural features influence individuals' interaction patterns and relationship formation (Kleinbaum, Stuart, and Tushman, 2013). Extensive research has

examined the role that individual differences play in predicting individual network position (Burt, Jannotta, and Mahoney, 1998; Mehra, Kilduff, and Brass, 2001). Though we find a short-term decrease in performance associated with one form of confirmation shopping, we find it intriguing that disconfirming feedback has the behavioral effect of triggering network changes that, in effect, increase recipients' likelihood of spanning boundaries or filling structural holes. At the organizational level, it's possible that distinct benefits will arise from a less constrained network. For example, less constrained individuals are more likely to birth creative innovations, yielding productivity benefits to the organization. Paradoxically, though, there is a tendency, particularly in relatively closed organizational systems, for network structures to, over time, become "closed". As described previously, over time, we tend to connect with those who are closely connected to our current set of connections, leading to incidences of "clusters" in the overall social network (Rivera, Soderstrom, and Uzzi, 2010). Perhaps one valuable, but unintended, side effect of disconfirming peer feedback is that it "upsets" this clustering tendency, leading, at least some actors in the network, to forge new, distant relationships that, though costly to performance in the short term, lead to innovative outcomes in the longer-term.

Limitations and Future Research

Our investigation has a number of strengths. First, by pairing archival field data with laboratory data, we document a practically significant phenomenon and confirm the hypothesized causal relationship and its mechanisms in the lab. Further, our field data are longitudinal, which allows us to observe employee behaviors over time. But our study also suffers from some limitations which serve as natural prompts for future research. Although we find that confirmation shopping is associated with lower performance in the subsequent year, we do not have data to show whether there are longer-term positive effects associated with the shifts in one's network in response to negative feedback. It is not lost on us that negative feedback

might ironically lead to more advantageous network positions for the recipient. (Burt, 2005, 2009) suggested that lower constraint is associated with a higher degree of social capital. It is possible that the longer-term effects include more collegial relationships with colleagues—that is, that employees reshape their networks until they find a better fit. Fit seems to be important, and the maintenance of a healthy self-concept is certainly important, at least in the longer term (Chatman, 1989; Elsbach, 2003). It is not inconceivable that the “shopping for confirmation” behavior could lead to longer-term positive performance. Future research should explore the long-term benefits associated with network shifts motivated by disconfirming feedback.

We believe that viewing feedback as a relational intervention provides an important alternate lens through which we can design performance improvement interventions; specifically, we can look to relationship science for insight into ways we might resolve these challenges. Relational closeness almost certainly matters here; we are unlikely to drop spousal or close friendship relationships that are the source of disconfirming feedback. But closeness is, in part, a reflection of the complex multi-dimensionality of a relationship. We don't drop spouses and close friends who provided disconfirming feedback, because those relationships, in some other way, are also a source of self-concept confirmation. That is, in some other dimension of the relationship, these close others separately provide us with the needed self-confirmation. One possible lesson from this research—one that we hope future work will explore—is that coupling developmental insight with other, separate, relational interactions that bolster the employees' self-concept may help to solidify relationships, and moderate the effects we observe in our data.

Our study does not explore the role of individual differences in the effects of negative feedback on network formation. We intuitively sense that individual differences matter; those who are inclined to seek feedback are very possibly less inclined to experience the feedback as a threat to the self-concept or to engage in confirmation-shopping behavior. How can we, then,

cultivate in employees the desire for developmental insight? Further, all feedback dimensions are not equal; it seems likely that certain dimensions of feedback are more meaningful, or central, to different employees, implying differential effects across different dimensions. We hope that future research will explore employee identification with feedback dimension as a moderator of our observed effects.

Finally, our field context is such that employees had the ability to actively reshape their network in response to feedback-activated threats to their self-concept. We hope that this study prompts an increased interest in understanding employees' coping mechanisms in organizational contexts where individuals do not have the ability to reform their network of relationships.

CONCLUSION

Feedback processes are nearly ubiquitous in modern organizations. Managers employ these processes naively, assuming employees will respond to them with dutiful efforts to improve. But we find that disconfirming feedback shakes the foundation of a core aspect of employees' self-concept, causing them to respond by reshaping their networks in order to shore up their professional identity and salvage their self-concept. This reshaping of employee networks contributes to lowered performance—a result ironically at odds with the ultimate goal of performance feedback. Our research offers an expanded view of social capital in interpersonal settings, and suggests that organizations must find ways to fulfill employees' need for a socially bolstered self-concept—that developmental feedback in the absence of this self-confirmation offers little hope for improving performance outcomes.

References

Adler, S., M. Campion, A. Colquitt, A. Grubb, K. Murphy, R. Ollander-Krane, and E. D. Pulakos
2016 "Getting rid of performance ratings: Genius or folly? A debate." *Industrial and Organizational Psychology*, 9: 219–252.

Aguinis, H., H. Joo, and R. K. Gottfredson
2011 "Why we hate performance management—And why we should love it." *Business Horizons*, 54: 503–507.

Aguinis, H., and C. A. Pierce
2008 "Enhancing the relevance of organizational behavior by embracing performance management research." *Journal of Organizational Behavior*, 29: 139–145.

Alicke, M. D., and C. Sedikides
2009 "Self-enhancement and self-protection: What they are and what they do." *European Review of Social Psychology*, 20: 1–48.

Ashford, S. J.
1986 "Feedback-seeking in individual adaptation: A resource perspective." *Academy of Management Journal*, 29: 465–487.

Ashford, S. J., R. Blatt, and D. V. Walle
2003 "Reflections on the looking glass: A review of research on feedback-seeking behavior in organizations." *Journal of Management*, 29: 773–799.

Baldwin, T. T., M. D. Bedell, and J. L. Johnson
1997 "The social fabric of a team-based MBA program: Network effects on student satisfaction and performance." *Academy of Management Journal*, 40: 1369–1397.

Borgatti, S. P., and P. C. Foster
2003 "The network paradigm in organizational research: A review and typology." *Journal of Management*, 29: 991–1013.

Brass, D. J.
1981 "Structural relationships, job characteristics, and worker satisfaction and performance." *Administrative Science Quarterly*, 331–348.

Brockner, J., and E. T. Higgins

2001 "Regulatory focus theory: Implications for the study of emotions at work." *Organizational Behavior and Human Decision Processes*, 86: 35–66.

Burt, R. S.

2005 *Brokerage and Closure: An Introduction to Social Capital*. Oxford University Press.

2009 *Structural Holes: The Social Structure of Competition*. Harvard university press.

Burt, R. S., J. E. Jannotta, and J. T. Mahoney

1998 "Personality correlates of structural holes." *Social Networks*, 20: 63–87.

Campbell, D. J., and C. Lee

1988 "Self-appraisal in performance evaluation: Development versus evaluation." *Academy of Management Review*, 13: 302–314.

Cannon, M. D., and R. Witherspoon

2005 "Actionable feedback: Unlocking the power of learning and performance improvement." *The Academy of Management Executive*, 19: 120–134.

Carver, C. S., and M. F. Scheier

1982 "Control theory: A useful conceptual framework for personality–social, clinical, and health psychology." *Psychological Bulletin*, 92: 111.

Chatman, J. A.

1989 "Improving interactional organizational research: A model of person-organization fit." *Academy of Management Review*, 14: 333–349.

Cleveland, J. N., K. R. Murphy, and R. E. Williams

1989 "Multiple uses of performance appraisal: Prevalence and correlates." *Journal of Applied Psychology*, 74: 130.

Crocker, J., K. Voelkl, M. Testa, and B. Major

1991 "Social stigma: The affective consequences of attributional ambiguity." *Journal of Personality and Social Psychology*, 60: 218.

Cross, R., and J. N. Cummings

2004 "Tie and network correlates of individual performance in knowledge-intensive work." *Academy of Management Journal*, 47: 928–937.

Culbertson, S. S., J. B. Henning, and S. C. Payne
2013 "Performance appraisal satisfaction." *Journal of Personnel Psychology*. Retrieved from <http://econtent.hogrefe.com.ezp-prod1.hul.harvard.edu/doi/full/10.1027/1866-5888/a000096>

Cummings, J. N., and R. Cross
2003 "Structural properties of work groups and their consequences for performance." *Social Networks*, 25: 197–210.

Davies, P. G., S. J. Spencer, and C. M. Steele
2005 "Clearing the air: identity safety moderates the effects of stereotype threat on women's leadership aspirations." *Journal of Personality and Social Psychology*, 88: 276.

Day, D. V., and L. M. Sulsky
1995 "Effects of frame-of-reference training and information configuration on memory organization and rating accuracy." *Journal of Applied Psychology*, 80: 158.

DeNisi, A., and C. E. Smith
2014 "Performance appraisal, performance management, and firm-level performance: a review, a proposed model, and new directions for future research." *Academy of Management Annals*, 8: 127–179.

Dutton, J. E., and E. D. Heaphy
2003 "The power of high-quality connections." *Positive Organizational Scholarship: Foundations of a New Discipline*, 3: 263–278.

Dutton, J. E., and B. R. E. Ragins
2007 "Exploring positive relationships at work: Building a theoretical and research foundation."

Dweck, C. S.
1986 "Motivational processes affecting learning." *American Psychologist*, 41: 1040.

Edmondson, A.
1999 "Psychological Safety and Learning Behavior in Work Teams." *Administrative Science Quarterly*, 44: 350–383.

Elsbach, K. D.

2003 "Relating physical environment to self-categorizations: Identity threat and affirmation in a non-territorial office space." *Administrative Science Quarterly*, 48: 622–654.

Fein, S., and S. J. Spencer

1997 "Prejudice as self-image maintenance: Affirming the self through derogating others." *Journal of Personality and Social Psychology*, 73: 31.

Ferris, G. R., T. P. Munyon, K. Basik, and M. R. Buckley

2008 "The performance evaluation context: Social, emotional, cognitive, political, and relationship components." *Human Resource Management Review*, 18: 146–163.

Gino, F.

2008 "Do we listen to advice just because we paid for it? The impact of advice cost on its use." *Organizational Behavior and Human Decision Processes*, 107: 234–245.

Greenwald, A. G.

1980 "The totalitarian ego: Fabrication and revision of personal history." *American Psychologist*, 35: 603.

Greve, H. R.

2003 *Organizational Learning from Performance Feedback: A Behavioral Perspective on Innovation and Change*. Cambridge University Press.

Grey, R. J., and D. Kipnis

1976 "Untangling the performance appraisal dilemma: The influence of perceived organizational context on evaluative processes." *Journal of Applied Psychology*, 61: 329.

Harris, M. M., and J. Schaubroeck

1988 "A meta-analysis of self-supervisor, self-peer, and peer-supervisor ratings." *Personnel Psychology*, 41: 43–62.

Healy, P., and G. Serafeim

2015 "An Analysis of Firms' Self-Reported Anticorruption Efforts." *The Accounting Review*.

Heneman, H. G.

1974 "Comparisons of self-and superior ratings of managerial performance." *Journal of Applied Psychology*, 59: 638.

Hofmann, D. A., M. A. Griffin, and M. B. Gavin

2000 "The application of hierarchical linear modeling to organizational research."

Retrieved from <http://doi.apa.org.ezp-prod1.hul.harvard.edu/psycinfo/2000-16936-011>

Ilgen, D., and C. Davis

2000 "Bearing bad news: Reactions to negative performance feedback." *Applied Psychology*, 49: 550–565.

Johnson, J. W., and K. L. Ferstl

1999 "The effects of interrater and self-other agreement on performance improvement following upward feedback." *Personnel Psychology*, 52: 271–303.

Jordan, A. H., and P. G. Audia

2012 "Self-enhancement and learning from performance feedback." *Academy of Management Review*, 37: 211–231.

Kadushin, C.

2002 "The motivational foundation of social networks." *Social Networks*, 24: 77–91.

Kay, E., H. H. Meyer, and j. R. P. French Jr.

1965 "Effects of threat in a performance appraisal interview." *Journal of Applied Psychology*, 49: 311.

Kleinbaum, A. M., T. E. Stuart, and M. L. Tushman

2013 "Discretion within constraint: Homophily and structure in a formal organization." *Organization Science*, 24: 1316–1336.

Kluger, A. N., and A. DeNisi

1996 "The effects of feedback interventions on performance: a historical review, a meta-analysis, and a preliminary feedback intervention theory." *Psychological Bulletin*, 119: 254.

Korsgaard, M. A.

1996 "The impact of self-appraisals on reactions to feedback from others: the role of self-enhancement and self-consistency concerns." *Journal of Organizational Behavior*, 17: 301–311.

Kunda, Z.

1987 "Motivated inference: Self-serving generation and evaluation of causal theories." *Journal of Personality and Social Psychology*, 53: 636.

Leary, M. R., and R. M. Kowalski

1997 *Social Anxiety*. Guilford Press.

Leary, M. R., E. S. Tambor, S. K. Terdal, and D. L. Downs

1995 "Self-esteem as an interpersonal monitor: The sociometer hypothesis." *Journal of Personality and Social Psychology*, 68: 518.

London, M., and J. W. Smither

1995 "Can multi-source feedback change perceptions of goal accomplishment, self-evaluations, and performance-related outcomes? Theory-based applications and directions for research." *Personnel Psychology*, 48: 803.

Ma, D.

2015 "Social belonging and economic action: Affection-based social circles in the creation of private entrepreneurship." *Social Forces*, 94: 87–114.

Marsden, P. V.

1987 "Core discussion networks of Americans." *American Sociological Review*, 122–131.

Mayhew, B. H.

1980 "Structuralism versus individualism: Part 1, shadowboxing in the dark." *Social Forces*, 59: 335–375.

Mehra, A., M. Kilduff, and D. J. Brass

2001 "The social networks of high and low self-monitors: Implications for workplace performance." *Administrative Science Quarterly*, 46: 121–146.

Meyer, H. H.

1980 "Self-Appraisal of Job Performance." *Personnel Psychology*, 33: 291–295.

Mohrman Jr, A. M., S. M. Resnick-West, E. E. Lawler III, M. J. Driver, M. A. Von Glinow, and J. B. Prince

1989 *Designing Performance Appraisal Systems: Aligning Appraisals and Organizational Realities*. Jossey-Bass.

Murphy, K. R.

1982 "Difficulties in the statistical control of halo." Retrieved from <http://psycnet.apa.org.ezp-prod1.hul.harvard.edu/journals/apl/67/2/161/>

Murphy, K. R., and J. Cleveland

1995 Understanding Performance Appraisal: Social, Organizational, and Goal-Based Perspectives. Sage.

Murphy, K. R., and J. N. Cleveland

1991 Performance Appraisal: An Organizational Perspective. Allyn & Bacon.

Murphy, K. R., J. N. Cleveland, and C. J. Mohler

2001 "Reliability, validity, and meaningfulness of multisource ratings." Handbook of Multisource Feedback, 130–148.

Murphy, K. R., and R. DeShon

2000 "Interrater correlations do not estimate the reliability of job performance ratings." Personnel Psychology, 53: 873–900.

Nag, R., K. G. Corley, and D. A. Gioia

2007 "The intersection of organizational identity, knowledge, and practice: Attempting strategic change via knowledge grafting." Academy of Management Journal, 50: 821–847.

Neville, L., and N. Roulin

2016 "Genius or Folly? It Depends on Whether Performance Ratings Survive the 'Psychological Immune System.'" Industrial and Organizational Psychology, 9: 281–288.

Ones, D. S., C. Viswesvaran, and F. L. Schmidt

2008 "No new terrain: Reliability and construct validity of job performance ratings." Industrial and Organizational Psychology, 1: 174–179.

Posthuma, R. A., and M. A. Campion

2008 "Twenty Best Practices for Just Employee Performance Reviews." Compensation & Benefits Review, 40: 47–55.

Pulakos, E. D., R. M. Hanson, S. Arad, and N. Moyer

2015 "Performance management can be fixed: An on-the-job experiential learning

approach for complex behavior change." *Industrial and Organizational Psychology*, 8: 51–76.

Pulakos, E. D., N. Schmitt, and C. Ostroff

1986 "A warning about the use of a standard deviation across dimensions within rates to measure halo." *Journal of Applied Psychology*, 71: 29.

Raudenbush, S. W., and A. S. Bryk

2002 *Hierarchical Linear Models: Applications and Data Analysis Methods Vol. 1*. Sage.

Raudenbush, S. W., A. S. Bryk, Y. F. Cheong, R. T. C. Jr, and M. du Toit

2016 *HLM7 Hierarchical Linear and Nonlinear Modeling User Manual 1 edition*. Skokie, IL: Scientific Software International, Incorporated.

Reagans, R., and B. McEvily

2003 "Network structure and knowledge transfer: The effects of cohesion and range." *Administrative Science Quarterly*, 48: 240–267.

Rivera, M. T., S. B. Soderstrom, and B. Uzzi

2010 "Dynamics of dyads in social networks: Assortative, relational, and proximity mechanisms." *Annual Review of Sociology*, 36: 91–115.

Roberson, L., S. Torkel, A. Korsgaard, D. Klein, M. Diddams, and M. Cayer

1993 "Self-appraisal and perceptions of the appraisal discussion: A field experiment." *Journal of Organizational Behavior*, 14: 129–142.

Ruehlman, L. S., S. G. West, and R. J. Pasahow

1985 "Depression and evaluative schemata." *Journal of Personality*, 53: 46–92.

Sedikides, C., and M. D. Alicke

2012 *Self-Enhancement and Self-Protection Motives*. Oxford handbook of motivation, ed. R. Ryan. Oxford University Press.[rWvH].

Sedikides, C., and M. J. Strube

1995 "The multiply motivated self." *Personality and Social Psychology Bulletin*, 21: 1330–1335.

Sherman, D. K., and G. L. Cohen

2006 "The psychology of self-defense: Self-affirmation theory." *Advances in Experimental Social Psychology*, 38: 183.

Sparrowe, R. T., R. C. Liden, S. J. Wayne, and M. L. Kraimer
2001 "Social networks and the performance of individuals and groups." *Academy of Management Journal*, 44: 316–325.

Swann Jr, W. B.
1983 "Self-verification: Bringing social reality into harmony with the self." *Social Psychological Perspectives on the Self*, 2: 33–66.

Swann, W. B., and S. J. Read
1981 "Self-verification processes: How we sustain our self-conceptions." *Journal of Experimental Social Psychology*, 17: 351–372.

Taylor, S. E., and J. D. Brown
1988 "Illusion and well-being: a social psychological perspective on mental health." *Psychological Bulletin*, 103: 193.

Tornow, W. W.
1993 "Perceptions or reality: is multi-perspective measurement a means or an end?" *Human Resource Management*, 32: 221–229.

Vaillant, G. E.
1995 *The Wisdom of the Ego*. Harvard University Press.

Viswesvaran, C., D. S. Ones, and F. L. Schmidt
1996 *Comparative Analysis of the Reliability of Job Performance Ratings*. American Psychological Association.

Wellman, B.
1997 "Structural analysis: From method and metaphor to theory and substance." *Contemporary Studies in Sociology*, 15: 19–61.

Wills, T. A.
1981 "Downward comparison principles in social psychology." *Psychological Bulletin*, 90: 245.

Wyer, R. S., and D. Frey

1983 "The effects of feedback about self and others on the recall and judgments of feedback-relevant information." *Journal of Experimental Social Psychology*, 19: 540-559.

Figures and Tables

Figure 1. Theoretical model

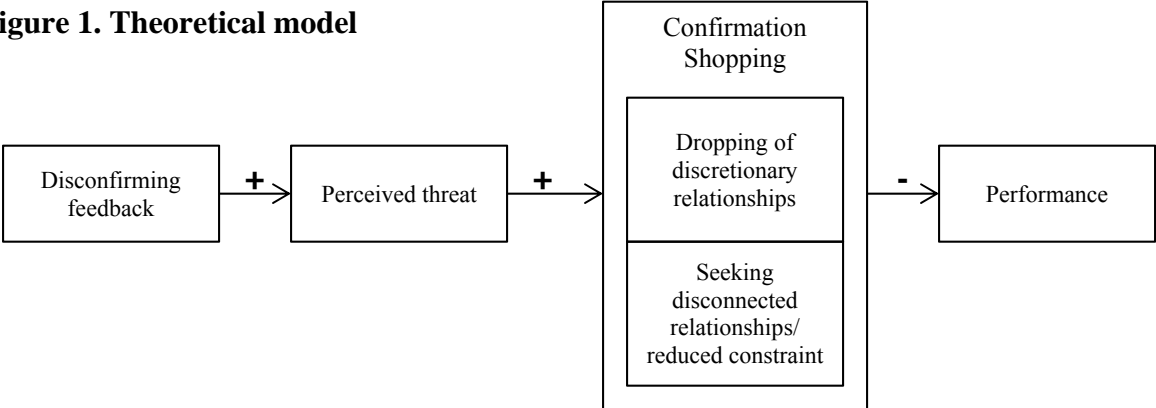


Figure 2: Breakdown of self and peer-evaluations, and employee relationships, by year

Year	Network Data (Captured in January)			Evaluation Data (Captured in December)		
	Unique Employees	Relationships	Avg. Relationships/Employee	Count of Self-evaluations	Count of Peer-evaluations	Total evaluations
2012	261	2172	8.32	229	1836	2065
2013	268	2368	8.84	252	1689	1941
2014	286	2335	8.16	261	1568	1830
2015	278	2185	7.86			

Figure 3: Timeline of data generation for performance analyses




Table 1: Individual level descriptive statistics and correlations

Variable	Mean	s.d.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1 Connections, 2012	11.09	6.43	1																								
2 Connections, 2013	11.55	6.05	.87	1																							
3 Connections, 2014	11.17	6.03	.80	.90	1																						
4 Connections, 2015	10.17	5.19	.72	.80	.91	1																					
5 Self-rating, 2012	4.96	.47	-.14	-.07	-.09	-.08	1																				
6 Self-rating, 2013	4.48	.52	-.12	-.20	-.18	-.11	.41	1																			
7 Self-rating, 2014	3.53	.44	-.13	-.13	-.12	-.08	.40	.40	1																		
8 Received rating, 2012	5.12	.37	.17	.17	.20	.21	.12	.17	.17	1																	
9 Received rating, 2013	4.67	.38	.21	.15	.20	.22	.01	.23	.19	.80	1																
10 Received rating, 2014	3.67	.33	.21	.24	.31	.33	.04	.30	.21	.65	.75	1															
11 Received rating, 2015	3.82	.40	.23	.15	.20	.26	-.01	.28	.15	.50	.64	.74	1														
12 Reviewer standardized rating, 2012*	.00	.05	.28	.23	.27	.28	-.04	.12	.08	.82	.72	.64	.48	1													
13 Reviewer standardized rating, 2013*	.00	.05	.28	.29	.30	.29	-.07	.11	.10	.78	.77	.62	.50	.79	1												
14 Reviewer standardized rating, 2014*	.00	.06	.26	.26	.31	.29	-.04	.14	.15	.57	.58	.75	.55	.60	.72	1											
15 Count of disconfirming reviews received, 2012	2.88	2.46	.22	.26	.15	.12	.64	.16	.15	-.28	-.29	-.22	-.15	-.28	-.31	-.26	1										
16 Count of disconfirming reviews received, 2013	2.57	2.77	-.02	.06	-.02	-.03	.27	.58	.19	-.12	-.29	-.09	-.10	-.13	-.14	-.09	.34	1									
17 Count of disconfirming reviews received, 2014	2.26	2.44	.07	.07	.09	.04	.20	.05	.62	-.21	-.23	-.29	-.27	-.20	-.17	-.18	.32	.28	1								
18 Obligatory disconfirming reviews received, 2012	1.16	1.45	-.06	.06	-.08	-.10	.51	.11	.15	-.27	-.29	-.24	-.27	-.25	-.31	-.25	.65	.25	.26	1							
19 Obligatory disconfirming reviews received, 2013	1.27	1.71	-.21	-.02	-.12	-.12	.26	.35	.11	-.22	-.35	-.20	-.34	-.24	-.23	-.21	.29	.67	.25	.52	1						
20 Obligatory disconfirming reviews received, 2014	1.16	1.48	-.04	.05	.02	-.03	.20	-.13	.43	-.27	-.33	-.39	-.39	-.32	-.28	-.22	.30	.14	.78	.49	.35	1					
21 Constraint, 2012	.22	.14	-.70	-.57	-.53	-.51	.10	.02	.02	-.18	-.25	-.28	-.35	-.34	-.25	-.19	-.17	.03	-.06	.12	.26	.13	1				
22 Constraint, 2013	.22	.14	-.70	-.74	-.67	-.61	.10	.06	.05	-.22	-.25	-.30	-.32	-.33	-.33	-.24	-.20	-.05	-.06	.06	.15	.11	.81	1			
23 Constraint, 2014	.22	.14	-.66	-.66	-.75	-.66	.10	.09	.08	-.22	-.25	-.26	-.30	-.30	-.30	-.23	-.09	.03	-.10	.21	.23	.10	.75	.85	1		
24 Constraint, 2015	.31	.21	.00	-.03	-.06	-.01	.00	-.02	.01	-.06	-.13	-.16	-.02	-.17	-.12	-.18	.13	.11	.10	-.11	.02	.00	-.06	.01	-.03	1	
25 Log bonus received, 2015	.09	.05	.22	.29	.36	.35	-.05	.03	-.04	.39	.34	.36	.31	.39	.44	.34	-.05	-.01	-.06	-.11	-.08	-.09	-.25	-.35	-.35	.01	1

n = 347 unique employees

Table 2: Effect of disconfirming feedback on the likelihood of an employee dropping the relationship in the subsequent year. Three-level logistic regressions, odds ratios reported.

VARIABLES	(1) Odds of Dropped Relationship Controls Only	(2) Odds of Dropped Relationship
Disconfirming Feedback (dummy)		1.367*** (0.165)
Same Unit (dummy)	0.639*** (0.0661)	0.636*** (0.0657)
Same Location (dummy)	0.816* (0.0952)	0.807* (0.0942)
Count, confirming reviews	0.983 (0.0193)	0.996 (0.0199)
Reviewer-standardized feedback score	0.492 (0.315)	0.930 (0.636)
Feedback giver, disconfirming review (dummy)	1.435*** (0.148)	1.503*** (0.158)
Gender (F=1)	1.000 (0.175)	0.978 (0.170)
Same gender (dummy)	1.252* (0.156)	1.252* (0.156)
Self-rating	0.904 (0.0863)	0.868 (0.0837)
Tenure (years)	0.955*** (0.0110)	0.954*** (0.0110)
Constant	0.458 (0.223)	0.462 (0.224)
Observations	3,982	3,982
Number of groups	310	310

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3: HLM results, predicting the effect of obligatory disconfirming reviews on change in network constraint.

VARIABLES	(1) Controls Only	(2) Year to Year Change in Constraint
Count, obligatory disconfirming reviews		-0.0169*** (0.00623)
Count, total disconfirming reviews	0.0105*** (0.00375)	0.0170*** (0.00444)
Count, total confirming reviews	0.00649** (0.00254)	0.00596** (0.00253)
Standardized average feedback score	-0.00525 (0.0110)	-0.0104 (0.0111)
Self-rating	-0.0180 (0.0251)	-0.0123 (0.0250)
Gender (F=1)	-0.0226 (0.0164)	-0.0207 (0.0163)
Constant	0.0702 (0.129)	0.0444 (0.129)
Observations	618	618
Number of groups	282	282

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 4: Confirmation shopping and subsequent year performance

VARIABLES	(1)	(2)	(3)	(4)
	H5a Log Bonus %	H5a Log Bonus %	H5b Log Bonus %	H5b Log Bonus %
	Controls Only		Controls Only	
Count, dropped 2014 disconfirming reviews		-0.00214		
		(0.00727)		
Total # of disconfirming reviews	-0.00334*	-0.00284		
	(0.00170)	(0.00207)		
H4 Residual				0.0320**
				(0.0125)
2015 Average peer-rating	0.0280**	0.0293***	0.0322**	0.0289**
	(0.0109)	(0.0103)	(0.0127)	(0.0111)
2014 Average peer-rating	0.0357	0.0351	0.0339	0.0379
	(0.0216)	(0.0209)	(0.0246)	(0.0226)
Gender (F=1)	-0.0291	-0.0284	-0.0291	-0.0307
	(0.0179)	(0.0178)	(0.0184)	(0.0189)
Constant	-0.142**	-0.145**	-0.159***	-0.161***
	(0.0564)	(0.0601)	(0.0543)	(0.0537)
Observations	132	132	132	132
R-squared	0.222	0.223	0.198	0.222
Business Unit Fixed Effects?	YES	YES	YES	YES
Number of Business Units	26	26	26	26

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5: Summary statistics by condition, Study 2.

Condition	Perceived threat	Anxiety after receiving feedback	Positive affect	Negative affect	Chose to change partner
Confirming feedback	1.54 (0.88)	1.34 (0.67)	3.04 (0.84)	1.25 (0.50)	8.8% (0.29)
Disconfirming feedback	2.67 (1.38)	1.35 (0.60)	2.67 (0.84)	1.32 (0.52)	29.7% (0.46)

Note. Standard deviations are reported in parentheses.

Appendix A: Description of residual regression methodology used in determining the relationship between confirmation shopping and subsequent year performance (testing Hypothesis 5b)

In determining the effect of confirmation shopping on subsequent year performance (our Hypothesis 5b) we employ a methodology presented by Healy and Serafeim (2015). A simple approach to testing our hypothesis would be to simply regress each employee's 2015 performance on their change in constraint from 2014 to 2015. Though we have conducted that analysis with results consistent with our Hypothesis 5b, our goal is to more precisely understand how differences in employees' network reshaping behavior in response to disconfirming feedback influence performance—not merely the relationship between changes in constraint and performance. Employees make a conscious decision as to how to respond to disconfirming feedback. Our analyses show that disconfirming feedback from obligatory colleagues is associated with decreases in network constraint in the subsequent network period. But all employees don't respond in precisely the same way to disconfirming feedback: some may be particularly extreme reshapers of their network, seeking new, extremely distant, connections. Others may show restraint, working to resist the tendency to reshape their network, and rather intensify their focus on developmental efforts, yielding no decrease (or perhaps even an increase) in constraint. In analyzing the effect of confirmation shopping on performance, we desire to understand the nature of these varying behavioral responses to instances of disconfirming feedback: do those who engage in confirmation shopping to an extreme degree (whom we deem "shoppers") perform more poorly in the subsequent year than those who exhibit restraint, and don't indulge the tendency to seek relatively disconnected colleagues?

Healy and Seraphim (2015) make the conceptual argument that regression residuals denote abnormally high or low instances of the outcome of interest, given the independent

variable and all relevant controls. Applied specifically to our context, the residuals of our test of Hypothesis 4 provide us with a distribution of constraint changes, as a function of the number of colleagues providing disconfirming feedback. More specifically, the residuals from our analysis of our Hypothesis 4 provide us with a distribution of employees. Some of these employees engaged to an extensive degree in confirmation shopping, yielding relatively extreme decreases in constraint given the number of disconfirming reviews received, in turn yielding a negative residual from our analysis of Hypothesis 4. Others (whom we refer to as “improvers” in our analysis) engaged in confirmation shopping to a lesser degree, yielding relatively muted decreases in constraint given the number of disconfirming reviews received, and in turn, yielding a positive residual from our analysis of Hypothesis 4.

In short, the residuals from Hypothesis 4 become a calculated variable reflecting the variance in change of constraint as a function of disconfirming feedback received. We interpret this variation in constraint as the variation in confirmation shopping behavior as a function of disconfirming feedback received. By employing these residuals in our analysis of Hypothesis 5b, we can predict the effect of relatively extreme, or muted, confirmation shopping behavior on subsequent year performance.