Influence flows in the academy: Using affiliation networks to assess peer effects among researchers

Craig M. Rawlings *, Daniel A. McFarland

Stanford University, 520 Galvez Mall, Rm. 510, Stanford, CA 94305-3084, USA

ARTICLE INFO

ABSTRACT

Little is known about how influence flows in the academy, because of inherent difficulties in collecting data on large samples of friendship and advice-seeking networks over time. We propose taking advantage of the relative abundance of "affiliation network" data to assess aggregate patterns of how individual and dyadic characteristics channel influence among researchers. We formulate and test our approach using new data on 2034 faculty members at Stanford University over a 15-year period, analyzing different affiliations as potential influence channels for changes in grant productivity. Results indicate that research productivity is more malleable to ongoing interpersonal influence processes than suggested in prior research: a strong, salient tie to a colleague in an authority position is most likely to transmit influence, and most forms of influence are likely to spill over to behaviors outside those jointly produced by collaborators. However, the genders and institutional locations of ego-alter pairs significantly affect how influence flows.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

Given the increasingly social nature of academic research, especially in the sciences (Adams et al., 2005; Powell et al., 2005; Wuchty et al., 2007), there is surprisingly little known about how academic colleagues influence one another's professional activities. Rather than examining interpersonal influences, prior research has tended to focus on the individual and institutional characteristics associated with different levels of productivity (Crane, 1965; Long, 1978, 1981, 1990, 1992; Long and McGinnis, 1981; Neumann, 1977; Reskin, 1978a,b; Xie and Shauman, 2003). This has given a somewhat static impression of the factors affecting research productivity; although peer influence has been suggested as a possible mechanism for certain aggregate findings (see Allison and Long, 1990; Azoulay et al., 2008; Bercovitz and Feldman, 2008). The lack of a more fully developed interpersonal explanation is somewhat surprising given the abundant research on the underlying social structures of scientific communities, which implies ongoing influence processes (Breiger, 1976; Friedkin, 1978, 1998; Moody, 2004; Newman, 2001). Research suggests that faculty members' professional contacts are important conduits for the evaluation and spread of specific research practices, and that status and gender are key moderating factors in these processes (Leahey, 2003, 2006). And yet, the general patterns of when "invisible colleges" are most likely to affect research practices remain unclear (Crane, 1972).

Our goal is therefore to begin to formulate and test a social network explanation for short-term changes in researchers' grant productivity, using a sample of faculty members who are regularly involved in grant activity at Stanford University over a 15-year period. Our setting exemplifies the trend toward large well-funded interdisciplinary research, and is therefore indicative of a growing number of research settings. It is therefore well-suited to address questions that are increasingly

---

important in understanding the social nature of research: do collaborators and colleagues influence one another's grant productivity? If so, are certain kinds of relationships more important in the spread of productivity norms and practices? Does influence emerge only through synergistic activity occurring in the immediate relationship, or is there evidence that influence spills over into other relationships and independent practices? Do social distances between individuals affect the ways that affiliations channel influence? Our approach offers preliminary answers to these questions, and may be used as a first evaluation of aggregate peer influence patterns in this population. These questions are not only theoretically important to academic researchers, but substantively important for administrators and funding agencies. If research productivity is shaped by ongoing social relations, this suggests that initiatives aimed at increasing collaboration and the collective awareness of productivity gains may be important strategies for increasing organizational productivity. And yet, there are significant obstacles in testing if and when such effects exist.

For many researchers, potential peer effects are threats to causal validity that must be made ignorable through experimental or quasi-experimental designs (Rubin, 1990); while for others, peer effects are analytically important but extremely difficult to quantify (Manski, 1993, 1995). In general, the difficulties in modeling peer influence have led researchers to focus on smaller groups, laboratory experiments and unique quasi-experiments (Azoulay et al., 2008; Christakis and Fowler, 2007; Friedkin, 1998). Many social network approaches are built upon the premise that peer influence is ubiquitous, but complete network information is prohibitively difficult to obtain on large samples over time. And yet, there are many times when researchers would like to know if evidence supports a peer influence explanation for important outcomes in a population and whether or not influence flows as one might expect based upon prior research. Are some types of relationships and types of individuals more influential than others? Given the rarity of natural experiments and the near impossibility of collecting even partially complete information on large peer networks over time, this question is largely unanswered for many important populations and processes.

In this article, we develop an approach that uses "affiliation networks" to assess large-scale patterns of peer influence. Affiliation networks exist whenever actors belong to the same groups (see Wasserman and Faust, 1994). Faculty members, for example, are professionally affiliated with one another through shared memberships in academic departments, dissertation committees, research teams, etc. Our approach assumes that a shared affiliation is a proxy for interpersonal contact, especially if the affiliation group is small and socially cohesive.[1] Affiliation networks tap into how interactions are organized around various "social foci" that transcend physical proximity, and are important in generating and sustaining social ties (Feld, 1981). Using affiliations therefore sacrifices precision in gauging the amount of peer influence involved in a given outcome, but opens research opportunities to assess peer influence in the aggregate and to test when certain individual and dyadic characteristics interact with different social foci in better channeling influence. We also suspect in general that affiliation networks will tend to underestimate peer-influence effects because they will often dilute the influence spread through contacts by including non-significant others.

Our models confront well-known obstacles in causal estimation (see Bramoullé et al., 2007; Hoxby and Weingarth, 2005; Mouw, 2006). There are numerous ways researchers may incorporate affiliation-based information into longitudinal models to gauge influence, and careful consideration should be given in selecting one's strategy. Ultimately, we follow a dyadic approach that is consistent with recent work using directed network data to gauge peer influence on health outcomes (Christakis and Fowler, 2007, 2008; Fowler and Christakis, 2008). We model effects between pairs of academic researchers who share a given affiliation – while incorporating individual-level fixed effects to create more conservative estimates than in prior studies. The limitations of our approach preclude precise assessments of the overall amount of peer influence, while affording *relative* comparisons of effects based upon different group affiliations and dyadic characteristics. As in recent work on health outcomes, we offer further validation of our models by examining certain asymmetries in dyads, because we would not expect effects to be contingent upon directionality if results are uniformly shaped by spurious shifts affecting productivity. Nonetheless, our approach is best suited for more exploratory analyses or for broad confirmatory studies, rather than as a precise estimate of peer influence within a given process.

As a first test of this approach we analyze the associations between affiliation networks and short-term changes in faculty grant productivity. Grant activity is an important and somewhat understudied area of faculty productivity, compared to publication rates. Data on grant activity is relatively available from central offices and therefore easier to collect and more accurate in comparison to publication data, which entails either exhaustive coding of faculty CVs or digitally "scraping" various online article indices.[2] Lag times on publications also vary widely from field to field in ways that could bias models, whereas grant applications and awards tend to have a more uniform annual periodicity. In general, we believe grant activity is a conservative test of our approach because grant applications and awards are less publicly known than faculty publications, and therefore less likely to influence productivity through broadcast effects rather than direct interpersonal contact. We take into account multiple affiliations in assessing: (1) how different relationships may serve as relatively more or less effective conduits

---

[1] As part of an attempt to validate this assumption, we have completed a survey of all current faculty members at Stanford (with a response rate of 49%). Results confirm that affiliations differ in how likely they are to indicate a subjective understanding of social contact and influence. Shared dissertation committee work is generally not considered a form of social contact, while joint publications are likely to be seen as contact (although this likelihood diminishes if a publication has a very large number of co-authors). More complete results from this survey, which validate our overall understanding of affiliation networks as proxies for social contact, are available from the authors upon request.

[2] At the time of writing, we are still incorporating publication data into our project through a highly involved set of online database scraping and author name disambiguation strategies.