NORTHWESTERN
UNIVERSITY

Guido Lorenzoni
Breen Family Professor of Economics

November 9, 2015

Dear recruiting committee,

I am writing regarding Ruben Gaetani, a PhD candidate from Northwestern on the job market this year. I know Ruben very well and I had the pleasure to advise him on all his projects. Ruben is an extremely energetic and passionate researcher, with a main interest in growth and technological innovation. Ruben is both interested in building theoretical models and in going to the data, both to uncover new stylized facts and to test his hypotheses. He has been remarkably entrepreneurial in putting together novel datasets to answer his questions and he has been a positive force in the department, in particular by working with other students and creating very fruitful collaborations.

Ruben's job market paper, coauthored with Enrico Berkes, is on the benefits of geographic concentration of innovative activity. A common hypothesis is that density promotes innovation, by facilitating interactions and formal and informal exchanges among researchers, entrepreneurs, and managers. In particular, this is often proposed as an important reason behind the economic success of some urban areas. Ruben and Enrico's first contribution is to put this idea to test, by constructing a dataset in which they associate patents to geographical locations and evaluate whether more densely populated areas produce a larger number of patents. Ruben and Enrico do a careful job of checking that the geographic information is reliable and fine enough for their purpose. They then discover that there is no clear pattern of correlation between density and patents, which, at a first pass, would seem to question the idea of density-driven innovation.

Looking more carefully at the data, Ruben and Enrico have noticed that, broadly speaking, there seems to be two types of innovations going on. First, some large firms have highly specialized labs,

PHONE: (646) 591-0581 EMAIL: guido.lorenzoni@northwestern.edu

often located in rural areas, near production facilities. These labs produce a lot of specialized innovation, which refines and improves existing technologies. Second, there are urban areas in which both large and smaller firms have labs that have a broader scope and that tend to produce less conventional innovations, i.e., innovations that can potentially start a whole new technology or product, or that set the stage for future applications. Their hypothesis is that the second type of innovations, which they label "unconventional innovations," are the ones that really benefit from population density, with the network of exchanges associated to it. To put this hypothesis to test they need to develop a measure of "unconventionality." To do so, they exploit the fact that they have access to data on cross citations among patents and they define and calculate a score that captures how much a patent combines ideas from disparate areas of knowledge. From a measurement point of view, this is probably the most innovative part of the paper. They then look at the relation between density and unconventional innovations and find that here there is indeed a positive relation, while no relation is present for conventional innovations. Of course, this evidence is essentially descriptive. So Ruben and Enrico also make several attempts at establishing a clearer causal link between density and unconventional innovations. In particular, they focus on what happens when a large firm decides to locate in a certain area and on how it impacts the pace of innovation at firms already present in the vicinity. Overall, it's probably fair to say that they do not have a completely exogenous source of variation to deliver foolproof evidence of a causal link. However, they present a very impressive body of evidence that seems certainly consistent with their hypothesis and that suggests the presence of a robust stylized fact to motivate future research.

The next contribution of the paper is to build and simulate a geographical model of innovation that embeds their basic hypothesis that density helps unconventional innovation. In the model, there are two specialized areas of research, labeled "design" and "software" and there is a continuum of locations. Researchers are permanently attached to a research area, but can choose the location where to operate. If a researcher specialized in one of the two areas, say design, comes up with a new idea, he can implement it in two ways: either combine it with an existing design idea and obtain a conventional innovation, or combine it with a new software idea and obtain an unconventional innovation. Incidentally, this modeling approach is consistent with the way in which they measure "unconventionality" in the data. Going back to the model, unconventional innovations can only happen in a location where both types of researchers are active. On the other hand, they assume that

being in a location heavily specialized in your line of research increases the probability of coming up with a new idea. This determines a trade-off in the choice of locations. Being surrounded by similar researchers increases your productivity, but lowers the chances of making unconventional innovations. In equilibrium researchers must be indifferent among all active locations. This indifference condition determines the distribution of researchers, the number of locations that specialize in only one area, the number of locations that specialize in both areas, and the population density in each location. The model predictions are consistent with the stylized facts documented in the first part of the paper, with more densely populated areas specializing in unconventional innovation and less densely populated areas specializing in conventional innovations. The model provides a useful setup to understand the forces at work and the authors also use it to take a first pass at understanding welfare implications and thus implications for policy. Even though they allow for a local planner to internalize externalities at the local level, the competitive equilibrium is still inefficient, due to the endogenous determination of locations and of their specialization pattern. The paper provides some interesting preliminary results on the benefits of policies that favor agglomeration of innovative activity.

Ruben also has a very interesting paper, joint with Matteo Li Bergolis, which investigates the relation between scientific discoveries and business investment at the micro level. Ruben and Matteo try to disentangle the effects of scientific discoveries, looking both at discoveries that eventually lead to profitable applications and at discoveries that fail to succeed commercially. They dub them, respectively, "successes" and "dead-ends". They construct a new dataset that connects patents developed by firms in Compustat with scientific papers appearing in the Web of Science. This allows them to trace specific industrial innovations to the scientific papers that provided the underlying fundamental research breakthrough, so they can identify "scientific shocks" at the root of R&D activity and physical investment. They build a model that captures the uncertainty associated to the early stages of implementation of a scientific breakthrough. In the model, firms solve a signal extraction problem, to find out the profitability of a given innovation they need to invest, that allows them to gradually learn the value of the innovation, and then decide whether to invest further or abandon it. The model predictions are then used to distinguish empirically successes and dead-ends in their sample, using ex post information on realized profitability and essentially using a type of sign-based restriction: both shocks lead to higher investment in the short run, successes lead to

higher profits and investment in the long run, dead-ends lead to a decrease in profits in the medium run (due to capital overaccumulation). The most interesting outcomes of this empirical analysis are first that they find large effects of scientific discoveries on investment activity and second that they show that the process by which the quality of the innovation is revealed is quite slow, leading to prolonged periods of over-investment in the case of dead-ends (of course, an over-investment that can only be identified ex-post).

Both the projects I described are very likely to produce follow-up work (some of which Ruben and coauthors have already started), given the range of questions that can be addressed with the datasets developed in these papers. Overall, my impression is that Ruben is a rare case of a graduate student who has already built a clear agenda and shown the skill and determination to pursue it in the long run. His willingness to combine a variety of data sources to answer his questions and his ability to move from data to models and back, show great promise for future work. Ruben is one of the two best students I have seen since moving to Northwestern three years ago (the other one is Arlene Wong, who is also on the market this year) and he has my strongest recommendation.

Yours truly,

Guido Lorenzoni