WEINBERG
COLLEGE OF
ARTS & SCIENCES

Department of Economics
2001 Sheridan Road, Room 302
Evanston, Illinois 60208-2600 USA

www.econ.northwestern.edu

Judd A. and Marjorie Weinberg
College of Arts and Sciences
Northwestern University

econ@northwestern.edu
Phone 847-491-5140
Fax 847-491-7001

NORTHWESTERN
UNIVERSITY

November 11, 2015

Recommendation for Ruben Gaetani

Dear Colleague:

I am writing to enthusiastically recommend Ruben Gaetani for a position in your department. Ruben specializes in macroeconomics, and I am his main advisor. His dissertation research is concerned with the determinants and consequences of innovation activity. In my view, Ruben is truly outstanding. He is the best student I have ever worked with, and (together with Arlene Wong, who is also on the market this year) the best macro student from Northwestern since I joined the department in 2008. He has developed an exciting research agenda entirely on his own, he is extremely hard working and ambitious, and has strong skills covering the entire range from original data work to theoretical modeling. I recommend him enthusiastically to all economics departments and business schools, and *especially the very best*.

Ruben's job market paper "The Geography of Unconventional Innovation" (joint with Enrico Berkes, a fourth-year graduate student at Northwestern) starts out as an empirical examination of one of the big ideas in economic growth (expressed in Lucas 1988, among many other places), namely that interpersonal spillovers fostered by density (i.e., people living and working closely together in dense cities) are a key driver of innovation. This idea is motivated by the observation that economic activity is concentrated in cities, and that increasing urbanization and increasing density is closely associated with economic growth at least since the industrial revolution. There is an existing empirical literature that has found evidence broadly consistent with a link from density to innovation, namely, metropolitan areas with a more dense population have a higher rate of patenting per capita. However, Ruben and Enrico point out that metropolitan areas are too large to provide a good test of a link from density to innovation: an MSA (metropolitan statistical area) generally includes large suburban and ex-urban areas in addition to an urban core.

Hence, the first contribution of their paper is to test the idea of a density-innovation link at a much more disaggregated level. Using the universe of patents filed in the U.S. between 2000 and 2010, they georeference each patent at the sub-county level. Sub-counties are much smaller than MSAs, for example, in New York City each borough is a separate sub-county. Using this data, they show that among sub-counties that have any innovation activity (i.e. at least one patent filed), there is no positive association between density and innovation; depending on controls, the relationship is either flat or decreasing.

This alone is a potentially important finding, and a less driven researcher may have stopped here (resulting in a publication in a solid field journal). However, Ruben and Enrico use this empirical

finding merely as a starting point for the main part of their analysis. They thought about potential reasons why the density-innovation link does not hold up in disaggregated data, and noticed an important feature: Firms with high rates of patenting generate much of their innovation in easily accessible suburbs or exurbs instead of the urban core (such as IBM in Armonk, NY or Motorola in Schaumburg, IL). If one thinks a bit about the intuition behind a density-innovation link, this makes a lot of sense. Density should matter for innovation if it relies on serendipitous exchanges between strangers that are fostered by interaction in close quarters. However, firms such as IBM and Motorola do not rely on serendipity for generating innovation: they employ researchers who work on specific, well identified problems related to their business, and the interactions between them are formalized inside the firm. Hence, density is less important, so that locating in suburbs with easy accessibility and cheap real estate makes sense.

Based on this intuition, Ruben and Enrico conjecture that density matters only for a specific kind of innovation, namely for innovation that is based on serendipitous exchanges of ideas between people from different backgrounds, who would be unlikely to interact within a given firm. The second important contribution of their paper is to find a way to actually measure this kind of innovation. They introduce the distinction between "conventional" and "unconventional" innovation, where unconventional innovation is an innovation that combines ideas from disparate fields of knowledge. They measure conventionality by analyzing the patterns of citation in a patent: an unconventional patent is one that cites other patents from an "unusual" (in a well-defined sense) combination of fields (this is inspired by the literature on citation patterns in scientific publications). For instance, a Motorola patent that only cites mobile-phone technology would be conventional, whereas a patent that cites mobile-phone technology alongside, say, textiles or sports equipment would be unconventional. They show that the conventional-versus-unconventional distinction has meaningful consequences: Unconventional patents are significantly more likely to have high impact (in the sense of being cited by future patents). Next, Ruben and Enrico establish the second main empirical finding of their paper: There is a strong positive association between population density and the rate of unconventional innovation.

Ruben and Enrico conjecture that a higher rate of serendipitous interactions in dense places is what explains the observed relationship. While the data does not provide a source of exogenous variation that would provide direct evidence of a causal link between density and unconventional innovation, they provide indirect evidence that supports such a link, in the spirit of a diff-in-diff exercise. Namely, they look at what happens to the patenting of established firms in a given location when a firm from a different field sets up a research establishment (i.e., an establishment that generates patents) in the same locality. They find that in the subsequent years, the existing firms are significantly more likely to file patents that combine their own area of expertise with that of the newly established firm (again based on patent citations). While there is more than one potential explanation for such a pattern, it is very suggestive of the density-driven mechanism that Ruben and Enrico have in mind.

In the final part of their paper, Ruben and Enrico develop a spatial growth model that can account for the facts that they establish in the empirical analysis. The basic framework is a model of vertical innovation, in which there is a variety of products, the quality of which can be improved by innovation over time. Ruben and Enrico extend this framework in two directions in order to speak to the facts that they established. First, innovation relies on two different fields of

knowledge, dubbed "software" and "design." Intuitively, one can think of a cell phone that needs both an operating system and a physical appearance (such as Apple's famous rounded edges), both of which can be patented. Analogous to the definitions in their empirical work, in the model a conventional innovation is an innovation that improves only one aspect of a given product, i.e., either software or design. Because such an innovation only has value when combined with an existing design, such innovations are sold to incumbents to improve their products. In contrast, an unconventional idea combines new ideas in both dimensions, resulting in an entirely new product that replaces the incumbent. This model of innovation is combined with a spatial model where people can locate in different cities. There are congestion externalities, but also positive spillovers that drive innovation. Most importantly, unconventional innovation relies on serendipitous encounters with innovators from a different field of knowledge, the likelihood of which increases with density. The model is kept tractable by including "city planners" who can use taxes and subsidies to internalize local externalities. Ruben and Enrico show that the model has a unique equilibrium that accounts for the facts they document in the data: some cities specialize in unconventional innovation and host a mix of innovators from different areas, whereas others cities are specialized in a specific type of conventional innovation. As in the data, the cities that generate unconventional innovation are denser that those that generate conventional ideas. The model also has important welfare consequences. Even though city planners internalize local externalities, the equilibrium is still inefficient due to global externalities generated by innovation. Ruben and Enrico show that policy interventions that push the urban structure closer to the social optimum can have large welfare consequences. Actual governments engage in a variety of policies that are aimed at affecting the spatial structure of innovation (i.e., public funding for "innovation incubators," general subsidies for R&D in specific places in attempts to create a new kind of "Silicon Valley," and also policies in support of existing firms or industries that aim to slow down structural change and corresponding changes in the urban structure), and Ruben and Enrico's paper is an important step in improving our understanding of how such policies work and whether they are desirable.

Ruben has a second completed paper (joint with Matteo Li Bergolis) on "The Economic Effects of Scientific Shocks." As is well known, much of the technological developments in recent decades have been spurred by specific scientific breakthroughs (such as technologies for data storage), and in other areas technology remains a bottleneck (such as energy storage). Ruben and Matteo construct a model of directed technological change that allows for scientific breakthroughs in particular areas. The twist is, however, that breakthroughs themselves are uncertain: breakthroughs take the form of "news" about future productivity which is subject to noise, so that what seems like a breakthrough does not always work out as expected. Ruben's model shows that not just the successes, but also the "dead ends" (i.e., the breakthroughs that in the end do not work out) have important implications for the path of technology, because the news of a future shock itself leads to additional investment in a sector, which is self-sustaining to a certain degree.

As the job market paper, this project has a fascinating empirical component. Ruben and Matteo combine data on patent citations and financial information on firms to empirically identify scientific breakthroughs that end up as successes or failures. The basic idea is that a breakthrough is any paper that ends up highly cited in patents. Success or failure is measured by the economic success of firms that strongly rely on a given innovation, as evidence by how often

the firm's own new patents cite the innovation. Hence, a failure would be a scientific innovation that leads to many cites (i.e., additional innovation), but where the firms that carry out that follow-up innovation end up being unprofitable. The end result of this procedure is a new data set on innovations for which there should be many uses (apart from applying it to the specific model described above).

Ruben has additional ongoing projects that are highly promising, and we are also working on a joint paper entitled "Employment Protection, Investment in Job-Specific Skills, and Inequality Trends in the United States and Europe." I will not describe this paper in detail so as not to praise what is partly my own work, except to say that the idea for this paper grew entirely out of Ruben's work (it is much closer to his research interests on innovation and technical change than to mine), and that working with him has been a wonderful experience.

One aspect of Ruben's personality as a researcher that I particularly like is that he sets very high standards for himself; he is self-critical almost to a fault, and if he realizes that some aspect of a project could have been done in a better way, he goes back and does everything again from scratch until he gets it right, regardless of the time and effort required. He is also a risk taker in the sense that he does not mind making major investments in a project before the payoff is clear. In the case of his job market paper, there was no way to know ex ante whether unconventionality would matter for patents' impact and whether it would end up being related to density in the way he conjectured. Not every project he started ended up having high payoffs, but he swings for the fences every time, and this is how I think important work gets done.

In summary, in my view Ruben is the kind of economist that every department should aim to hire. He is a remarkably mature researcher who sets his own agenda and combines innovative empirical work with theoretical modeling to advance the research frontier. He has already completed two papers that should have a good chance of getting published in a top journal, and I expect many more from him in the years to come. Ruben is also a great person to be around, and he has contributed to the macro group in various ways in the last few years. I think he has had a very positive impact on other students here (not least his coauthors), and he will contribute greatly to the group that hires him. I give him my highest recommendation, and I strongly urge you to consider him for an appointment in your department.

Please feel free to contact me if you have any additional questions.

Sincerely,

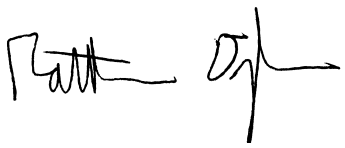

Matthias Doepke
Professor of Economics