

Discussion of “Venture Capital and the Transformation of Private R&D for Agriculture” *

MICHAEL EWENS

Caltech & NBER

December 15, 2019

Abstract

The authors document changes to the finance ecosystem for entrepreneurial firms in the agriculture sector. Using a unique combination of several proprietary databases, the empirical analysis reveals strong correlations between the flows to startups and signals of investment opportunities. The paper could be improved on three dimensions. First, a careful review of the database creation with particular attention paid to the variation in industry and coverage differences by data provider. Next, more evidence is needed to convincingly demonstrate that the financing environment changed in agriculture in ways different from that experienced in all of the startup ecosystem. Finally, the authors have many opportunities to explore how venture capital investments dynamics are connected to changes in the agriculture industry or changes to the supply side of the market.

*Ewens: mewens@caltech.edu, California Institute of Technology. 1200 East California Blvd. Pasadena, CA 91125. Phone: 619-512-3820.

1 Summary

This paper investigates the trends in global private equity (PE) and venture capital (VC) investments in agriculture startups. After finding evidence for rapid increases in capital flow to these startups – particularly after the financial crisis of 2008 – the paper explores the sources of changes. Specifically, the capital provided to these startups is growing relative to the supply of capital invested by the public sector and public firms. Next, regression analysis confirms that investments in agricultural sectors is strongly correlated with both past liquidity events in the sector and changes in prices for major commodities. The results are consistent with investors responding to investment opportunity signals in the agricultural space.

Answers to these questions are important for researchers and policy makers who aim to support the agricultural sector and its innovation. More broadly, the analysis of changing investor behaviors such as new allocations to new industries reveals where startup financing constraints lie. Finally, the results also contribute to perennial debates around public versus private R&D spending.

The comments on this paper focus on several topics including data construction, interpretation and suggestions for additional analyses.

1.1 Combining databases: benefits and pitfalls

The paper describes a major data exercise merging three databases:

[W]e construct a dataset of 4,500 startups in agriculture with more than 10,000 financial transactions, including investments and their exits, combined from three proprietary data sources: Crunchbase, PitchBook (Morningstar), and VentureSource (DowJones).

The authors should be applauded for combine these related, but distinct sources of data. However, such merges face challenges when data providers differ in their coverage and industry classification methodologies. Consider first the VentureSource database provided by Dow Jones. In my experience using this data, I have learned that their best coverage is for U.S.-based startups backed by venture capital. Informal conversations with the data provider also revealed that the data quality is high only after 1990 (the firm was founded in 1987). Next, Pitchbook provides significantly wider coverage by region than VentureSource. It was founded in 2007 by a former employee of VentureSource. Its early focus was on the U.S. PE ecosystem, but has grown – and in my opinion – improved over the last 5-7 years. Given their founding in 2007, it is not clear where their historical data was sourced, which is an important uncertainty when merging with databases that have longer time series. Finally, Crunchbase fills out the financing coverage. This database

is best known for its coverage of U.S. early-stage (often pre-VC) financings. It began in 2007 by a technology blogger to keep track of startups covered on its site and grew as a wiki-style site. It also appears to have benefited from the switch away from PDF to XML formatted regulatory filings for private equity exemption notices in 2009. Given its short history and narrow industry focus in the early years, its quality for agriculture startups is unclear.

My suggestion to the authors is to first motivate the merge of these three databases. For example, is there evidence that one has poor coverage of exits or non-software companies? The main concern is that the quality, coverage and definitions differ widely across (and possibly within) data providers. Note that each of these companies likely make most of their revenue from non-academic customers, which likely means that they are less concerned with historical data and have resources devoted to the current period. The best motivation of this merge would be to fill in gaps in each databases' coverage. Alternatively, the authors could pick one as a "master" dataset and use the remaining two to fill in coverage gaps or missing values.

As an example of possible time-varying coverage, consider a query of Pitchbook for U.S.-based agricultural startups financed in 1990–2018 for "all" investor types in Figure 1.


Several questions emerge from this figure that warrant some discussion in the paper. First, note the increase in counts that coincides with the founding of Pitchbook in 2007. Can the authors confirm that this is a real phenomenon rather than a coverage change? Second, for the paper's analysis of VC-backing changes over the sample period, how does one explain the red line of percent VC-backed? It exhibits significant variation, but no major break in the 2008-2009 period.

Next, I was able to conduct a quick merge of VentureSource and Pitchbook for U.S.-based VC-backed agricultural startups. The first challenge is selecting industries. Unlikely publicly-traded firms that have at least one SIC code associated with their SEC filings, private firms self-assign industries or data providers have their own methodology. Private firm data providers do not always use SIC or NAICS codes, which limits merging across databases. My attempt to get ag-startups in VentureSource and Pitchbook was thus challenging. For VentureSource, I choose the broad category "Agriculture and Forestry" with added flag in business descriptions for "farm," "harvest," and "agriculture." The latter is important because tractor guidance software categorized as "Software" but the startups have "farm" in their description. I followed the paper's approach to querying Pitchbook.

The merge of VentureSource and Pitchbook was done using the startup's name (after some basic cleaning). Figure 2 presents the counts of successful merge using the "Agriculture and Forestry" search (51 matches) and that combined with the keywords (105). The results indicate that casting a wide net in any search is important for increasing the sample size, but also that individual matches demand random checks for false positives.

Figure 1: Agriculture-tech startups: Pitchbook

The figure reports the count of VC/PE-backed U.S.-based ag-tech startups from 1990 to 2018. The left axis and blue bars report the raw counts of the number of startups raising such capital. The red line reports the fraction that are “VC-backed” according to Pitchbook.



A closer analysis of the failed merges suggests that some hand-collection is necessary. Aerial Intelligence (Figure 3) is found in Pitchbook, but not as an agriculture firm:

The startup’s description strongly suggests it is ag-tech. Next, the startup VinSense is not found in VentureSource. It has the following business description:

Developer of a crop management software designed to enhance crop uniformity and increase crop volume. The company’s software helps to improve crop management using soil sensors and offers optimal soil nutrient management, enabling producers, field managers and winemakers to manage soil moisture, pruning, irrigation, canopy management and water conservation.

An analysis of VinSense’s financings shows that it over half of its capital raised was in the form of government grants, while its equity investors made abnormally small investments. This example makes clear that data providers also have different methodologies on what constitutes “VC-backed.” It also shows how valuable merging different databases can be for improving cov-

Figure 2: Merge of VentureSource and Pitchbook

The figure reports the counts of successful merges of startups in VentureSource and Pitchbook using two different search queries for VentureSource.




Figure 3: Example startup in Pitchbook

An example description of a startup in Pitchbook not listed as “ag-tech.”

General Information

Description

Developer of agricultural data platform designed to offer insights and analytics for the agricultural commodities supply chain. The company's agricultural data platform is an application that converts big-data from satellites and weather stations into actionable, macro-level insights on agricultural production, enabling agricultural businesses to avail information of flooding, drought, and other adverse conditions that might affect their commodities.

Most Recent Financing Status (as of 13-Aug-2018)

The company raised \$100,000 of convertible debt financing from Mindset Ventures on February 20, 2017.

Website	www.aerialintel.com	Universe	Venture Capital
Entity Types	Private Company	Primary Industry	Media and Information Services (B2B)
Legal Name	Aerial Intelligence, Inc.	Other Industries	Database Software
Business Status	Generating Revenue		Business/Productivity Software
Ownership Status	Privately Held (backing)		Artificial Intelligence & Machine Learning
Financing Status	Venture Capital-Backed	Verticals	Big Data
			SaaS
			TMT

erage. I would like to see more discussion of the rules each data providers uses when classifying and collecting data for the industries of interest.


Because data providers change methodology and quality over time, merging them can create spurious trends. This leads to two concrete suggestions:

1. Run the analysis on each database separately and demonstrate results all go in the same direction.
2. Manual merge all three using a wide net on each database and document why they disagree.

1.2 What's the counterfactual?

The paper argues that there has been a meaningful increase in capital provided to agriculture startups. The authors could do more to tease out the overall or macro trends from these changes. For example, Figure 4 shows the number of startups backed by VC since 1990 (according to VentureSource). The vertical red line shows the break in financing proposed by the authors. Clearly, the overall VC market experienced a change around 2008. Thus, I suggest that the authors isolate an area *within* agriculture that grew differently. For example, one could conduct a structural break test with unknown break in mean (constant), repeating the exercise for the changes in agricultural prices.

Figure 4: Number of startups financed by VC



1.3 What are the next steps?

The paper has the opportunity to explore deeper issues in both agriculture and venture capital. The paper hints at one direction:

Several factors may have affected the hurdle rate [VC’s willingness to invest], such as an increase in the ratio of agricultural prices to non-agricultural commodity prices, the occurrence of large exit events in highly visible ag technologies, the emergence of new technological opportunities based on advances in enabling technologies (such as cheaper genome sequencing, genome editing, or data capacity of sensors and networks), as well as changes in (agricultural) labor markets both in high income and middle income countries.

The authors also highlight changes to input prices that could provide exogenous variation in empirical tests:


In agriculture, increases on minimum wage for agricultural workers may have an impact on agricultural production and prices, and potentially on the returns to investments in startups in the industry

One suggestion is to follow a similar strategy found in Ewens, Nanda and Rhodes-Kropf (2018). They study the impact of changing startup costs after the introduction of the cloud. Their focus was on the information technology sector, however, it is likely that impacts are likely *within* agriculture. The same technological shock could be used to study the role of investor value-add to this industry or how capital flows between different sectors in the agricultural space.

Another avenue for additional analysis begins with the premise that the increase in capital to agriculture is real. One can ask who are the investors? This is an interesting question because agriculture is a non-traditional space for both VC and private equity. One prediction is that existing investors are “pivoting” or moving into agriculture to exploit new technology in the space. Here, changes in investing do not represent a demand shock, but rather a spillover from lack of investment opportunities elsewhere. Alternatively, the new investments have new investor entrants that are VCs. Such a pattern is consistent with a supply-side shock or exits from established ag-tech firms. One way to investigate this issue would require tracking the work histories of the partners in the startup financings. Finally, it is possible that the growth in ag-tech is facilitated by new *types* of investors (e.g. incubators, corporate venture capital, PE hybrids). If so, then the facts would be consistent with existing VCs having real constraints on either skill or capital for agricultural investment opportunities. In fact, Figure 3 in the paper (recreated below) shows that

venture capital is not the only source of capital in these startups and there is some evidence of changing composition of investors.

Figure 5: Figure 3 from Silva, Graff and Zilberman (2019)



2 Conclusion

The paper documents changes to the entrepreneurial finance ecosystem in agriculture. Moreover, it documents strong correlations between the flows to startups and signals of investment opportunities. My suggestions for the authors are three-fold. First, a careful review of the database creation with particular attention paid to the variation in industry and coverage differences by data provider. Next, more evidence is needed to convincingly demonstrate that the financing environment changed in agriculture in ways different from that experienced in all of the startup ecosystem. Finally, the authors have many opportunities to explore how venture capital investments dynamics are connected to changes in the agriculture industry or changes to the supply-side of the market.