

What Can Gold Farmers Teach Us About Criminal Networks?

By observing how covert financial networks operate in online games like *World of Warcraft*, we can learn about how they might function offline.

By Brian Keegan, Muhammad Aurangzeb Ahmad, Dmitri Williams, Jaideep Srivastava, Noshir Contractor

DOI: 10.1145/1925041.1925043

As information communication technologies have grown more pervasive, understanding the ways in which they are employed for unexpected, deviant, or even criminal purposes becomes increasingly important. Networks in particular are growing in importance as a lens onto how individuals behave and organizations operate. It seems appropriate that network analysis methods could also be applied to understand how clandestine or criminal organizations, such as drug traffickers and terrorist cells, coordinate and adapt while avoiding detection. Yet, the hidden nature of the relationships in these dark networks necessarily implies that collecting and analyzing complete or even representative data on these networks is very difficult.

However, the explosion of behavioral data available in online databases has opened up new avenues of social research. One such source are so-called massively multiplayer online games (MMOGs), large-scale social environments in which players of varying levels of expertise join cooperative teams to accomplish complex tasks. To the extent that individuals in online virtual worlds engage in similar psychological, social, and economic behavior as they do in the offline world, research in MMOGs and virtual worlds can potentially be mapped backwards to illuminate offline behavior [1]. Because the organizations that operate MMOGs maintain archival databases of all player actions and attributes, it is

possible to analyze data on a scale that would be impracticable or impossible to do in the real world. For example, administering surveys to thousands of

“Gold farming and real-money trading refer to practices that involve the sale of virtual in-game resources for real money via exchanges outside the game.”

people is hugely expensive and peoples' cognitive biases and social desirability factors often makes results unreliable. Likewise, observing individuals interacting may be more reliable, but observations could start after or stop before interesting interactions take place.

MMOGs such as *World of Warcraft*, *EverQuest II*, and *Lord of the Rings Online* are examples of game worlds where millions of players interact in a persistent virtual environment. Just as these game economies exhibit macroeconomic characteristics observed in real-world economies [2], virtual worlds also contain black markets for acquiring goods and skills. The organization and behavioral dynamics of users engaged in these illicit operations

coupled with the exhaustive digital footprints of their activity potentially provides a unique window into understanding how clandestine networks operate in other contexts.

WHAT IS GOLD FARMING?

While playing alone or with other players, MMOG players accumulate experience, armor, spells, and weapons to improve their power against computer non-player characters (NPCs). The virtual goods and in-game currency players acquire make their characters more powerful, and the acquisition of these items is one of the major goals of play. Virtual goods like in-game currency, scarce commodities, and powerful weapons require substantial investments of time to accumulate. However, these can also be obtained from other players within the game through trade and exchange.

“Gold farming” and “real money trading” refer to practices that involve the sale of virtual in-game resources for real-world money via exchanges outside of the game. The name stems from a variety of repetitive routines (“farming”), which are employed to accumulate virtual wealth (“gold”), which is sold to other players who lack the time or desire to accumulate their own in-game capital [3]. By repeatedly killing NPCs and looting the currency they carry, farmers accumulate currency, experience, or other forms of virtual capital that they exchange with other players for real money via transactions outside of the game. Other types of activities also fall under the banner of “real money trade”; players can sell rare weapons, armor, and spells within the game for offline money and accounts with elite characters have been known to sell for thousands of dollars.

Gold buyers then consume the gold they purchased in the game to obtain more powerful weapons, armor, and abilities for their avatars. This, in turn, accelerates players to higher levels and allows them to explore larger parts of the game world, confront more interesting and challenging enemies, and increase their social standing without having to invest much time in the tedious parts of completing quests and killing monsters for money and items.

“Gold farming has been constructed as a deviant activity by both the game developers as well as the player communities.”

Gold farming and real money trade operations originated in the first massively multiplayer online role-playing game, *Ultima Online*, in 1997. The practice grew rapidly with the parallel development of an e-commerce infrastructure in the late 1990s and the introduction of MMOGs into East Asian markets. Gold farming operations now appear to be concentrated in China where the combination of high-speed Internet penetration and low labor costs has facilitated the development of the trade.

Although outwardly innocuous, gold farming has been constructed as a deviant activity by both the game developers as well as the player communities for a variety of reasons. First, in-game economies are designed with carefully calibrated activities and products that serve as sinks to remove money from circulation. The injection of farmed gold into the game economy creates inflationary pressure, unintended arbitrage opportunities, and other perverse incentives that undermine the stability of the game economy. Second, farmers’ activities often overtly affect other players’ experiences, for instance taking over profitable regions of the game and preventing other players from completing quests. Farmers also employ anti-social computer scripts (“bots”) to automate the farming process that results in the uncanny experience of zones filled with players but bereft of social interactions. Third, the game developers are risk-averse to the legal implications of sanctioning a multinational industry estimated to generate between \$100 million and \$1 billion in revenue annu-

ally [4] while lacking legal jurisdiction, precedent, or regulation.

Finally, farming upsets the meritocratic and fantasy-based nature of the game. Some players may cease to play if other players can buy rather than earn accomplishments. If you team up with someone clad in rare Titansteel armor, it might be an unwelcome surprise in the middle of battle if this outwardly expert player does not even know basic game mechanics. For these reasons, many of the largest game developers like Blizzard, Sony, and Square actively and publicly ban accounts engaged in gold farming. However, other games such as *Second Life*, *Project Entropia*, and *EVE Online* explicitly allow players to indirectly or even directly exchange online currency and items for offline money.

HOW DO GOLD FARMERS BEHAVE?

Our initial work approached gold farming as a binary classification problem: you are either a gold farmer or you are not. What behavioral attributes influenced whether or not accounts and their characters were identified as gold farmers? First, a training set of known farmers caught by human players and customer service representatives from Sony Online Entertainment (SOE) gave us our reference set. Then, using deductive approaches like logistic regression as well as inductive machine learning techniques, we examined an anonymized database constructed from both surveys and behavior data collected by the game maker. By combining self-reported demographics (like age, language, and gender) with gold farming behavior (like time played, money earned, NPCs killed), latent behavioral proxies (like quests completed, recipes learned, deaths), and behavioral patterns (like successive NPC kills), we classified players based on their likelihood of matching “caught” farmer’s patterns [5].

As expected, some variables such as speaking Chinese, playing for long periods of time, using recently established accounts, dying repeatedly, and avoiding quests greatly increased the odds of being identified as a gold farmer. This meshed well with journalistic and community accounts of gold farm-

ers being primarily located in China, playing accounts continuously until they are banned, not being skilled at non-routine in-game tasks, and avoiding the story-based elements of the game.

However, some variables which should have been obvious validity checks for whether an account was or was not involved in gold farming acted unexpectedly. Killing many NPCs actually significantly reduced the odds of being identified as a gold farmer while having a large amount of wealth in a bank did not significantly increase or decrease the odds. Farmers' behavior evidently did not differ significantly and was likely concealed by the behavior of other players. Certainly, elite and hardcore players have killed many NPCs and accumulated great wealth. Alternatively, farming accounts that have escaped detection also exhibit this pattern. Using machine learning classifiers to predict gold farming status, the precision and recall of these models was surprisingly low. In fact, the classifier algorithms were returning many false positives—accounts fitting the profile for other gold farmers but never identified and banned by the game administrators. Clearly our reference set was not an ideal set.

SOE administrators, like many law enforcement organizations, rely on reports from players, patrols through the game world, coordinated sting operations, and database sleuthing to identify and “roll up” gold farmers. Based on the administrators' experience, they observed farmers employing increasingly sophisticated organizations and supply chains they described like a drug trafficking operation. There are the “farmers” who actually collect the gold from the game environment and send into the distribution network, the “mules” who move this money between other agents in the network, the “dealers” who interact with customers to give them the virtual items, the short-lived “marketers” who spam chat channels, and the “wholesalers” or “bankers” who receive and distribute goods but otherwise remain inactive to avoid attention. Clearly, this was not a simple binary classification task—each of these roles had very distinct behavioral profiles.

HOW DO GOLD FARMERS INTERACT?

In light of the fact there are simultaneously many gold farmer roles or classes as well as a significant number of unidentified farming accounts, we shifted the attention of our analysis [6]. Because trade of items and in-game currency are the fundamental operations in a gold farming network, we decided to employ a network analytic perspective. The trade network was constructed from the list of the exchanges of items and money between players by direct interaction or in-game mail. These could include legitimate bartering (trading an item for an item), market exchange (trading an item for currency), as well as unreciprocated “gifts” (sending currency or item, but receiving nothing in return).

We examined the known farmer's prior trade exchanges with other characters that were never banned. This naïve suspect-by-association heuristic allowed us to define a set of farming *affiliates*: these could be the paying customers of a gold farmer, unidentified farmers, or unsuspecting players. The majority of players who never interacted with a farmer were classified as *non-affiliates*. This trade network had a number of interesting topological properties. In keeping with the mapping philosophy of comparing and testing these in-game networks to their real-world counterparts, we examined these features both for the game data and for a dataset of known Canadian drug traffickers developed by Carlo Morselli—known as the CAV-IAR network. Given that criminal and

clandestine organizations are assembled on underlying trust relationships, network analyses of trust proxies—like trade and exchange relationships—could reveal important patterns about how individuals in these organizations are distributed throughout the entire network.

First, the distribution of trade interactions was extremely unevenly distributed throughout the network—the network follows a power law degree distribution. The vast majority of characters in *EverQuest II*'s network have traded with one or two other characters over the whole nine-month span of data. However, there were several dozen characters out of the approximately 43,000 characters on the server in the trade network that traded with hundreds of other characters.

This is not a smooth power law either; rather it is a *truncated power law*. Were we to extrapolate a trend line in **Figure 1** for the number of nodes having between 1 and 10 connections, we see that beginning around 20 or 30 connections the data begins to fall off below this trend line. This pattern has been observed in other studies of social networks, transportation networks, and biological networks and indicates that at a transition point, it becomes much more expensive or difficult to make additional links. As a result, significantly fewer nodes are able or willing to make that extra connection. For example, it might not be much more difficult for most people to keep in regular touch with 50 members of their social network versus 40, but it may become substantially more difficult to remember names, backgrounds, and shared experiences after you have 100. This is not the case for everyone—a few people are still social butterflies who can keep up 500 or 1,000 connections—but there are substantially fewer of these butterflies than we would expect by simple power law extrapolation.

For both the suspect affiliates as well as the legitimate non-affiliates, this transition point occurs around the same point in the network; between 20 and 30 trade partners. However, for the gold farmers, the transition point occurs much earlier; between 5 and 8 trade partners. Paradoxically, this result suggests that although farmers

“As MMOGs grow in scale and complexity, these exciting worlds and the novel player interactions within them assume increasing social relevance.”

Figure 1: The number of trading links in the network is unevenly distributed such that a few characters have most of the connections. While it is difficult for typical accounts to maintain more than 20 or 30 connections, gold farmers actually tend to trade with fewer players than typical players.

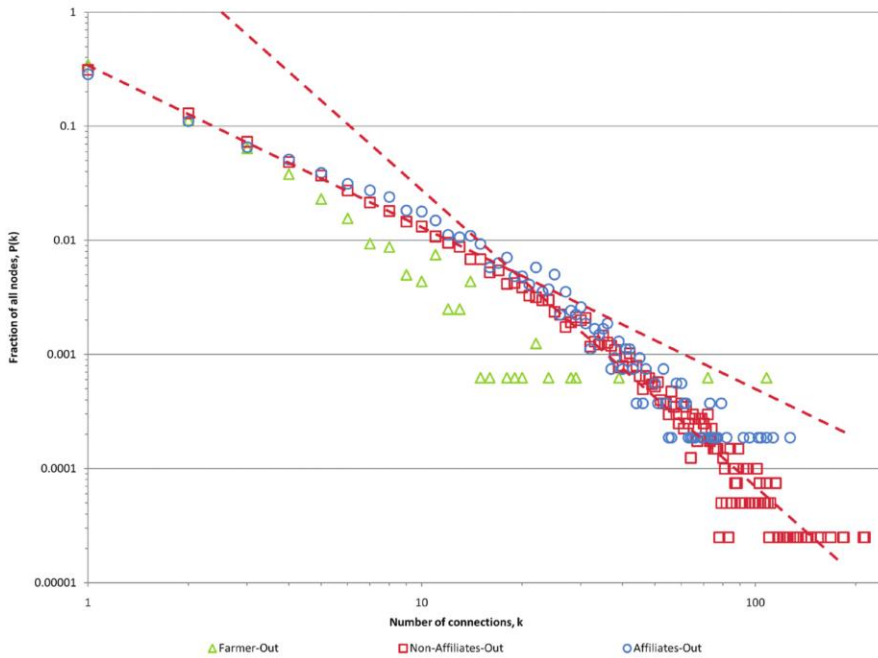
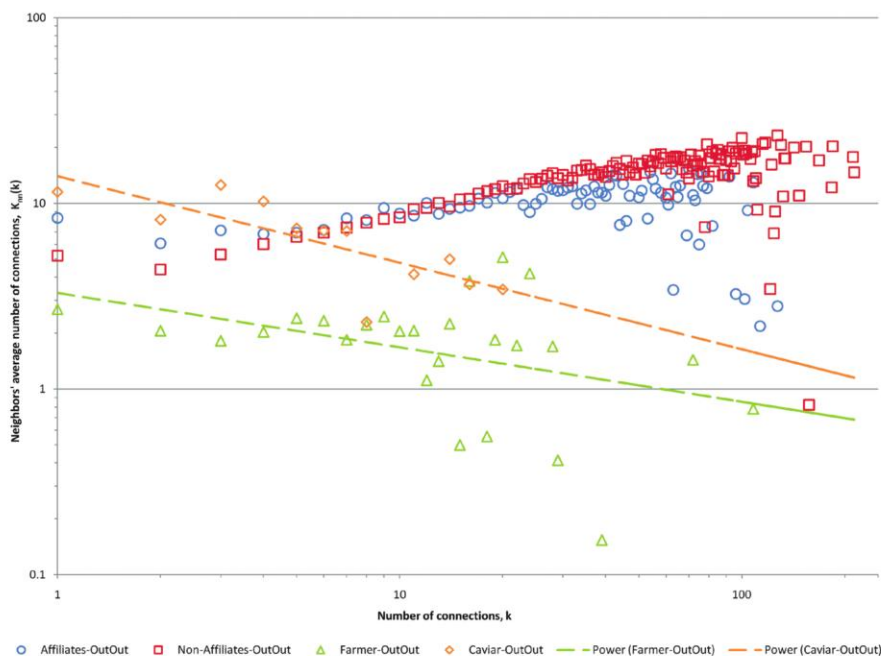


Figure 2: Typical players and unidentified affiliates with many connections trade with other characters who have many other connections. Gold farmers and offline drug traffickers with many connections prefer to trade with actors who have few other connections.



attempt to limit who they trade with, this effort is futile because they are still identified. Conversely, the affiliate accounts appear to avoid detection by emulating the behavior of the population at large. However agency in this context goes both ways; game administrators appear to only identify particular subclass of farmers with limited connectivity while many likely farmers not only continue to operate undetected, but remain well-connected.

Second, the three broad character categories we identified—farmer, affiliate, non-affiliate—occupy very different positions in the network relative to each other. Following our findings here, although farmers have fewer trade connections to other characters compared to non-affiliate characters, what few trade connections they have are employed much more intensely than we see among non-affiliate characters. Unlike farmers, affiliate characters have both more connections than non-affiliates but, like farmers, affiliates use these connections much more intensely than non-affiliates. This suggests that farmers and affiliates are both repeatedly exchanging gold and items with other trusted members of their organization. This point is further corroborated by examining differences in the tendency for characters' to form clustered trading patterns where characters A and C trade because they both also trade with B. Farmers have significantly higher clustering than non-affiliates, while affiliates have insignificantly less clustering. Farming characters form tight-knit trading groups with other trusted co-offenders, which must support important efficiencies because it is a distinct and intuitive signature for administrators to identify.

The third and final topological feature of the gold farming network we examined was its assortativity. We have an intuitive notion (certainly honed in middle and high school) that, rather than being randomly mixed together, well-connected people tend to be connected to other well-connected people while poorly connected people tend to be connected to other poorly-connected people. Statistical physicists like Mark Newman have termed this structural tendency as *assortativ-*

ity. The opposite pattern (*dissortativity*) of well-connected nodes preferentially connecting to poorly connected nodes and vice versa has been found in a number of ecological, biological, and technological networks. In other words, while assortativity is observed in networks characterized by sociality and collaboration, dissortative patterns appear to emerge in networks evolving in a context that demands resilience under selection pressures. It may be easier to attack a system where all the well-connected nodes are connected to each other than a system where well-connected nodes are insulated from each other through poorly connected nodes.

The concepts of assortativity and dissortativity also highlight a tension that exists in clandestine networks like gold farming trade networks [7]: should collaboration and efficiency be prized despite the heightened risks of detection or should resilience trump efficiency and flexibility? Our analysis (see **Figure 2**) demonstrates a very notable distinction. Affiliates and non-affiliates nodes clearly have assortative mixing patterns in which characters with many trading partners have a tendency for their neighbors to also be well-connected. However, the identified farmers exhibit a clear dissortative pattern. The CAVIAR network also exhibits a strong dissortative pattern and provides evidence of the mapping of online behaviors back to offline behavior; clandestine networks in both online and offline trafficking contexts preferred to insulate well-connected individuals from each other. Furthermore, a number of individual affiliate accounts are well outside of the typical range and are approximated much better by the gold farmer trend line. This suggests assortativity metrics may be useful for identification or prediction of not only gold farmers, but other well-connected members of clandestine organizations.

WHAT THIS MEANS FOR CRIMINAL ACTIVITY IRL

As MMOGs grow in scale and complexity, these exciting worlds and the novel player interactions within them assume increasing social relevance. In

“MMOGs potentially provide a window into our understanding of offline behavior that is otherwise difficult or impossible to study.”

particular, gold farming offers a case in which people are willing to pay a premium to enhance their experience in their leisure activities. However, the production of virtual wealth outside of community norms and rules offers a striking example of how illicit goods, clandestine organizations, and law enforcement demands and limitations also impinge on these game worlds. Given these parallels, MMOGs potentially provide a window into understanding of offline behavior that is otherwise difficult or impossible to study.

Obviously the notion of markets, property, and regulation within MMOGs raise a variety of complicated questions. What rights, if any, do players have over the digital artifacts they create and develop? What are the offline social, cultural, and technological contexts in which virtual economic production occurs? How should communities in online worlds respond to potentially exploitative phenomenon like inflation, arbitrage, and monopolies? How aggressively should administrators pursue deviants and traffickers to balance community stability with player privacy?

Very similar questions have preoccupied social scientists, philosophers, and political leaders since well before the advent of the Internet. Indeed, it is reassuring that despite how digitized and distributed these virtual worlds may be, MMOGs are still very human systems that evince the same struggles and debates that have preoccupied societies for ages. While the consequences of players' trials and tribulations

within MMOGs may not directly affect the “real world,” these online worlds are still very real sites for social interaction, organization, and economic trade. In that respect, MMOGs provide truly exciting platforms to both understand how human behavior unfolds using immaculately documented data as well as how to design technology to support and enhance social and economic interactions for both the online and offline world.

Biographies

Brian Keegan is a Ph.D. student at Northwestern University's School of Communication. He studies team assembly mechanisms under boundary conditions of online organization.

Muhammad Aurangzeb Ahmad is a Ph.D. student in Computer Science and Engineering at the University of Minnesota. He studies computational trust in virtual worlds and online reputation systems.

Dmitri Williams is an associate professor at the USC Annenberg School for Communication. He studies the social practices and implications of online communities and video games.

Jaideep Srivastava is a professor in computer science and Engineering at the University of Minnesota. He uses web mining to study online databases and multimedia systems to model human interactions.

Noshir Contractor is the Jane & William White Professor of Behavioral Sciences at Northwestern University. He studies the factors that lead to the formation, maintenance, and dissolution of dynamically linked social and knowledge networks in communities.

References

1. Williams, D. The mapping principle and a research framework for virtual worlds. *Communication Theory*, 20(2010), 451-470.
2. Castronova, E., Williams, D., Shen, C., Ratan, R., Xiong, L., Huang, Y. and Keegan, B. As real as real? Macroeconomic behavior in a large-scale virtual world. *New Media & Society* 11, 5 (2009), 685.
3. Heeks, R. *Current Analysis and Future Research Agenda on "Gold Farming": Real World Production in Developing Countries for the Virtual Economies of Online Games*. Institute for Development Policy and Management, University of Manchester, 2008.
4. Lehtiniemi, T. How big is the RMT market anyway. *Virtual Economy Research Network*, [Mar. 2 2007].
5. Ahmad, M. A., Keegan, B., Srivastava, J., Williams, D. and Contractor, N. Mining for Gold Farmers: Automatic Detection of Deviant Players in MMOGs, 2009.
6. Keegan, B., Ahmad, M., Williams, D., Srivastava, J. and Contractor, N. *Dark Gold: Statistical Properties of Clandestine Networks in Massively Multiplayer Online Games*. IEEE, 2010.
7. Morselli, C., Giguère, C. and Petit, K. The efficiency/security trade-off in criminal networks. *Social Networks* 29, 1 (2007), 143-153.