Multidimensional Platform Design[†]

By André Veiga, E. Glen Weyl, and Alexander White*

The defining feature of a platform business is its reliance on its own user base to produce value. The platform literature has extensively analyzed the incentives for subsidies this creates and the need for platforms to reach a "critical mass." Yet it is often insufficient for a platform to simply attract "enough" users. Instead, it needs the *right* users: those that will be profitable and/or attractive to other profitable users. In this paper, we analyze the static design and dynamic development of platforms to attract a valuable user base.

The literature on platforms (or "networks" as they were initially called) first gained prominence in the 1980s by studying products whose quality is determined by the total number of users (Katz and Shapiro 1985). In the last decade and a half, research on platforms has increasingly focused on "multi-sided markets" in which different, clearly distinguishable groups of users (e.g., game developers and players, reader and advertiser, shoppers and vendors) play distinctly different roles, with each group or "side" serving as a magnet to attract the other side (e.g., Rochet and Tirole 2003).

This work usually assumes that firms can charge a separate price to each group, which is false in many applications. Facebook and LinkedIn cannot maintain separate explicitly targeted offerings for popular users; newspapers do not explicitly offer different products to wealthy readers despite advertisers' greater willingness to pay to reach them. Consequently, platforms have an incentive to design non-price features to attract these valuable users, directly—and indirectly by appealing to other users whom valuable users find attractive. Moreover, summarizing a platform's quality requires a statistic richer than merely the *number* of users (as in Chandra and Collard-Wexler 2009).

In the next section, we propose a simple model of this phenomenon in a monopolistic one-sided platform/network by combining a platform model (Weyl 2010) with one of sorting by non-price product characteristics in selection markets (Veiga and Weyl 2016). In Section II, we analyze this model to characterize optimal platform design and dynamic deployment of a platform in terms of a notion of centrality proposed by Bonacich (1987). In Section III, we apply these ideas to help interpret the evolution of successful social networks and cities. Section IV concludes. Given space constraints, all technical details not necessary to follow the basic arguments appear in the online Appendix.

Veiga: Department of Economics and Nuffield College, University of Oxford, Manor Road Building, Oxford OX1 3UQ, UK (e-mail: andre.veiga@economics.ox.ac.uk). Weyl: Microsoft Research, 641 Avenue of the Americas, New York, NY 10011 and Yale University Department of Economics and Law School (e-mail: glenweyl@microsoft.com); White: School of Economics and Management and National Institute for Fiscal Studies, Tsinghua University, Weilun Building, 100084 Beijing China (e-mail: awhite@sem.tsinghua.edu.cn). This paper replaces and supersedes the now-defunct papers "Multi-Sided Platforms with Heterogeneous Externalities," "Platforms with Heterogeneous Externalities" and "Multidimensional Product Design," all by the first two authors, that have not been published already. We thank Ben Golub and Jonathan Levin for helpful comments and Chang "Leo" Liu for excellent research assistance. All errors are ours.

[†] Go to https://doi.org/10.1257/aer.p20171044 to visit the article page for additional materials and author disclosure statement(s).

I. Model

A monopolistic platform chooses a vector of platform characteristics ρ and a special distinguished characteristic $p \in \mathbb{R}$. The variable p is a "vertical" characteristic that is viewed as harmful by all users and is always beneficial to the platform; it may represent price but need not, as many platforms are offered free of charge. There is a unit mass of potential users each characterized by a type vector θ distributed with density f. The vector θ determines both the contribution that users make to the platform's value and their taste for using the platform. To model the first aspect, let $\phi(\theta)$ be a vector of user characteristics that impact both the platform's profit and its perceived value to users at large.

We assume that the profits of the platform and its value to users depend only on the aggregate value of these features. That is, if the set of users participating in the platform is Θ, then platform profits and user utilities are functions of ρ and $\Phi \equiv \int_\Theta \phi(\theta) f(\theta) \, d\theta$. We therefore refer to ρ as the platform's *exogenous quality* and to Φ as its *endogenous quality*.

User utility depends on both forms of quality as well as on the particular user's type. While we leave utility very general (subject to smoothness assumptions specified in the online Appendix), we make one notable assumption: users participate on the platform exactly if $u(\rho, \Phi; \theta) \geq p$. Given the substantial freedom available in choosing the units of u, this is largely a normalization; the substance of this assumption is that, since p is viewed as harmful by all users, regardless of any other features, raising p sufficiently high will always induce exit by any given user while lowering it sufficiently will induce entry. Finally, the platform makes profits $\pi(\rho, \Phi, p)$ that are strictly increasing in p everywhere.

As this model is somewhat abstract, consider two concrete special cases. First, suppose there is no ρ, ϕ is unidimensional and equal to one regardless of θ, and π takes the form $\Phi p - C(\Phi)$ for some scalar function C. This is a model of a one-sided platform monopolist whose costs and value proposition to users depend only on the *number* of other users, as in (Weyl 2010). Second, consider the case in which ρ is unidimensional, ϕ is two-dimensional with ϕ_1 equal to one regardless of θ and ϕ_2 is flexible, u depends only on ρ, θ, and p, but not on Φ, and π takes the form $\Phi_1 p - \Phi_2$. This is a model of product design in a market with adverse or advantageous selection, where ρ represents product quality, Φ_1 the quantity, and ϕ_2 the cost of a given customer, as in Veiga and Weyl (2016).

II. Design and Development

We now derive first-order conditions for the optimal design of a platform in steady state and a heuristic the platform can follow to induce this optimal design. A critical concept in our analysis is the sorting effect of a quality dimension. Let q denote an exogenous or endogenous quality dimension. We define

$$\sigma_{q\phi_i} \equiv M\text{cov}\left[\frac{\partial u}{\partial q}, \phi_i \Big| \partial\Theta\right]$$

as the *sorting effect* of q (given fixed values of all quality variables). Here, cov is the covariance operator, $\partial\Theta$ is the set of *marginal users* with $u = p$, and M is the density of users in this marginal set defined by $M = -\dfrac{\partial \int_\Theta f(\theta) \, d\theta}{\partial p}$. This measures the extent to which q is especially attractive to users with high values of ϕ_i and thus the extent to which increasing q will tend to differentially attract high ϕ_i users. Sorting effects are critical to determining the value of users in this environment, as they govern a platform's ability to selectively attract desirable users. We denote the matrix of sorting effects across dimensions of endogenous quality by $\Sigma_{\phi\phi}$, with typical element $\sigma_{\phi_i\phi_j}$.

A. *Design*

We then use this matrix to define the shadow value λ associated with endogenous quality according to the following matrix equation, which we then interpret as the central analytical idea of this paper:

$$(1) \quad \lambda \equiv \underbrace{\left(I - \Sigma_{\phi\phi}\right)^{-1}}_{\text{centrality term}}$$
$$\times \underbrace{\left(\frac{\partial\pi}{\partial\Phi} + \frac{\partial\pi}{\partial p} E\left[\frac{\partial u}{\partial\Phi}\Big|\partial\Theta\right]\right)}_{\text{direct effect}},$$

where I is the identity matrix. This somewhat abstract expression can be seen as an application of a notion of "centrality" in a graph/network first proposed by Bonacich (1987) and applied to economics by Ballester, Calvo-Armengol, and Zenou (2006). We now take a brief detour to introduce this notion and thus explicate equation (1).

Bonacich's notion of centrality applies to a graph where each node has associated with it some direct importance, and each directed edge between a pair of nodes represents the influence one node has on another. Bonacich centrality captures the notion that the value of every node is its direct value plus the sum over all nodes it influences, weighted by the influence it has on those nodes, of the value of those nodes. Thus, central nodes are those that are directly valuable and that influence nodes that are directly valuable and so on. If V is the vector of direct values and F is the matrix of directed influence, then Bonacich centrality is given by $(I - F)^{-1} V$.

Now consider the graph formed by the endogenous quality dimensions. The direct value of each node is given by the second factor on the right-hand side of equation (1) while the matrix of influences is given by Σ containing sorting effects of each dimension on the others. Given these identifications, the right-hand side of equation (1) is simply Bonacich's centrality measure in our setting.

To understand why the appropriate direct value is $\frac{\partial \pi}{\partial \Phi} + \frac{\partial \pi}{\partial p} E\left[\frac{\partial u}{\partial \Phi}\middle| \partial\Theta\right]$, it is useful to consider Spence's (1975) analysis of optimal monopolistic product quality and Weyl's (2010) generalization thereof to settings with multidimensional heterogeneity. In Spence's analysis, product quality has two effects: it has some direct impact on cost and it allows the firm to raise prices in proportion to the average marginal utility the quality creates for *marginal* consumers. These two effects correspond to the two terms in this direct value expression. However, in our setting, in addition to reflecting cost, the first term may capture the degree to which users of different types may be monetized by the platform through various channels such as advertising. The second term captures the degree to which users embody quality for which either the platform can directly charge (if p represents an actual price) or that it can exploit in order to cut its investment in exogenous quality.

To see why $\Sigma_{\phi\phi}$ is the appropriate matrix of influence, note that endogenous quality serves two roles: it directly allows the platform to profit from or charge for this quality, but, as we argued above, it also allows the platform to selectively attract users who contribute to endogenous platform quality of various kinds. This feedback loop amplifies (or dampens if the effects are negative) the impact of acquiring endogenous quality. This loop is captured by Bonacich's centrality metric.

Using these shadow values of endogenous quality, we can now define the optimal choice of the exogenous quality variables that the platform controls directly.

PROPOSITION 1: *Necessary conditions for the platform to choose an optimal setting of p and ρ are*

$$(2) \qquad \frac{\partial \pi}{\partial p} = M\lambda^\top E[\phi|\partial\Theta],$$

$$(3) \qquad -\frac{\partial \pi}{\partial \rho} = \frac{\partial \pi}{\partial p} E\left[\frac{\partial u}{\partial \rho}\middle| \partial\Theta\right] + \Sigma_{\rho\phi}\lambda,$$

where $\Sigma_{\rho\phi}$ has typical entry $\sigma_{\rho_i \phi_j}$.

Equation (2) dictates that the platform sets the direct marginal contribution to profits of p equal to its indirect cost of repelling M users who, on average, bring endogenous quality $E[\phi|\partial\Theta]$, which is valued at opportunity costs given by λ. Meanwhile, (3) implies the platform sets the marginal direct cost of exogenous quality, $-\frac{\partial \pi}{\partial \rho}$, equal to the sum of (i) the value the platform receives from the generalized "Spence term," and (ii) the value, in terms of *endogenous* quality, resulting from sorting of users, caused by the shift in exogenous quality.

B. *Development*

While Proposition 1 characterizes the optimal "rest point" for a platform, there is a well-known problem that platforms can face in trying to induce this. This is known as the "chicken and egg" problem (Caillaud and Jullien 2003), because (under the standard setup) users' utility depends on the number of users that participate, there may be multiple equilibria. For example, at a given price, pessimistic expectations about others' behavior may lead no users to participate. On the other hand, optimistic expectations can be self-confirming. How can a platform curtail this potential "failure to launch" (Evans and Schmalensee 2010)?

Building on earlier work on contingent pricing (e.g., Dybvig and Spatt 1983), Weyl (2010) proposed a strategy, the "insulating tariff," that platforms could use to ensure equilibrium uniqueness, in the game played by users, in the setting where only their preferences but not their costs/contributions are heterogeneous. If the platform charges a price that varies as a function of the number of users who actually choose to use the platform, it may ensure that the desired number of users participate, regardless of expectations. Insulating tariffs make the unique equilibrium the one desired by the platform. One way to interpret this strategy is in dynamic terms: the platform charges a low price (or even offers a subsidy) early on, when there are few users, gradually raising the price as the platform develops.

In our extended setting also featuring heterogeneous contributions and costs, the chicken and egg problem is richer. Because the value created by users is not a scalar but instead a vector of aggregate characteristics, Φ, there is a broader range of ways in which a platform might fail to launch. It could, for instance, attract a large mass of users but a sub-optimal mix of them. The set of equilibria could thus be very complicated. A policy of insulation is also trickier because the platform needs to be able to adjust its exogenous quality vector so as not only to ensure the right number of users regardless of their expectations but also the appropriate composition.

This has two implications for the feasibility of insulation. First, the dimensionality of exogenous quality (including p) relative to endogenous quality is crucial to determining whether insulation is possible. If endogenous quality is of higher dimension than exogenous quality, insulation will typically be impossible. If exogenous quality is of higher dimension than endogenous quality, there will often be many ways to insulate, at least locally. If the dimensions are the same then (at least locally) the insulating design will be unique.

Second, even when the dimensions align, insulation may not be possible because the exogenous quality dimensions may not be "strong enough" to counteract the effects of endogenous quality. A crucial feature of price that allows for insulation is that there is always a price high enough that no one is willing to pay it and a subsidy great enough that it can attract even the most skeptical users. Even if exogenous quality is of high dimension, it may not span a sufficient range of user utility to allow for insulation. In Section 3.1 of our online Appendix, we provide examples illustrating these issues.

However, even when full "insulation" is impossible, the platform may be able to use contingent pricing to guarantee, albeit less robustly, a unique equilibrium among users featuring its desired level of endogenous quality. Exogenous quality can insulate some key dimensions of endogenous quality, which can, in turn, reinforce other dimensions of endogenous quality, potentially causing a cascade that eventually sustains platform growth. A relevant example is White's (2013) model of a search engine, which offers unpaid "algorithmic" search results that attract searchers, who then attract advertisers. In this vein, also see Hagiu and Spulber (2013).

Section 3.2 of the online Appendix provides a stylized example of a social network where the multidimensional structure of exogenous quality and multiple layers of such a reinforcement strategy are used. We discuss the ideas behind this application informally in Section IIIB.

III. Applications

How might the above help explain certain observed strategies of cities and social networks?

A. *The Creative Class and Smart Cities*

A theme in contemporary urban theory (Florida 2002) has been the importance to cities of attracting the so-called "creative class," whose spillovers attract entrepreneurs and business. An extensive literature in economics has attempted to quantify these effects (e.g., Shapiro 2006), and many cities have adopted policies aimed at attracting these creative types (Markusen and Gadwa 2010). Typical policies involve subsidizing artist exhibitions/studio space and public purchases of local art.

Typical justifications of these policies, such as those offered by Florida (2002), are based not on claims of the direct benefits of artistic activity, but rather on the differential capacity of creative residents' presence to attract directly productive, high-taxpaying denizens. This causal chain is closely related to the logic of Proposition 1. Perhaps our analysis above, with appropriately specified heterogeneous preferences and contributions to amenities could build on the work of, e.g., Diamond (2016) to serve as an urban policy design framework.

B. *Rolling Out a Social Network*

Unlike most cities, social networks like Facebook and LinkedIn emerged over the last decade and a half. Given these platforms' youth, their operators have focused less on short-run profits and more on building durable foundations.

This includes carefully orchestrating the sequence of different kinds of users' arrival. For example, Kilpatrick (2011) recounts the story of Facebook's founder Mark Zuckerberg deciding to first target Harvard students, then other Ivy League students, and then, gradually, others, in seemingly concentric circles of increasing size

and diminishing social prestige. While "exclusivity" plays no role in Facebook's apparent steady-state strategy, this aspect appears to have been vital along the path to its current popularity. Our model provides an interesting way to interpret the course that Facebook followed in updating its features (including, e.g., its privacy policies) over time, as it became a mass market platform. Moreover, the (quasi-)insulating strategies discussed in Section IIB may serve as a useful approach to modeling this kind of situation, including using Bonacich centrality to identify which users to pursue when.

IV. Conclusion

This paper offers a model and some basic insights into the design and deployment of a monopoly platform in which (i) users make heterogeneous contributions to platform success, and (ii) they are attracted by the *characteristics* of other platform users rather than just their number. We hope that future research provides a more complete picture of dynamic platform strategies in environments with these traits. We also hope that researchers will analyze platform competition in this setting. In analyzing analogous selection markets, Veiga and Weyl (2016) find that competition can give rise to socially harmful "cream-skimming," which may be of policy concern for the regulators of platforms.

REFERENCES

Ballester, Coralio, Antoni Calvo-Armengol, and Yves Zenou. 2006. "Who's Who in Networks: Wanted: The Key Player." *Econometrica* 74 (5): 1403–17.

Bonacich, Phillip. 1987. "Power and Centrality: A Family of Measures." *American Journal of Sociology* 92 (5): 1170–82.

Caillaud, Bernard, and Bruno Jullien. 2003. "Chicken & Egg: Competition among Intermediation Service Providers." *RAND Journal of Economics* 34 (2): 309–28.

Chandra, Ambarish, and Allan Collard-Wexler. 2009. "Mergers in Two-Sided Markets: An Application to the Canadian Newspaper Industry." *Journal of Economics and Management Strategy* 18 (4): 1045–70.

Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000." *American Economic Review* 106 (3): 479–524.

Dybvig, Philip H., and Chester S. Spatt. 1983. "Adoption Externalities as Public Goods." *Journal of Public Economics* 20 (2): 231–47.

Evans, David S., and Richard Schmalensee. 2010. "Failure to Launch: Critical Mass in Platform Businesses." *Review of Network Economics* 9 (4).

Florida, Richard. 2002. *The Rise of the Creative Class: And How It's Transforming Work, Leisure, Community, and Everyday Life.* New York: Basic Books.

Hagiu, Andrei, and Daniel Spulber. 2013. "First-Party Content and Coordination in Two-Sided Markets." *Management Science* 59 (4): 933–49.

Katz, Michael L., and Carl Shapiro. 1985. "Network Externalities, Competition, and Compatibility." *American Economic Review* 75 (3): 424–40.

Kilpatrick, David. 2011. *The Facebook Effect: The Inside Story of the Company That Is Connecting the World.* New York: Simon & Schuster.

Markusen, Ann, and Anne Gadwa. 2010. "Arts and Culture in Urban or Regional Planning: A Review and Research Agenda." *Journal of Planning and Educational Research* 29 (3): 379–91.

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990–1029.

Shapiro, Jesse M. 2006. "Smart Cities: Quality of Life, Productivity, and the Growth Effects of Human Capital." *Review of Economics and Statistics* 88 (2): 324–35.

Spence, A. Michael. 1975. "Monopoly, Quality, and Regulation." *Bell Journal of Economics* 6 (2): 417–29.

Veiga, Andre, and E. Glen Weyl. 2016. "Product Design in Selection Markets." *Quarterly Journal of Economics* 131 (2): 1007–56.

Weyl, E. Glen. 2010. "A Price Theory of Multisided Platforms." *American Economic Review* 100 (4): 1642–72.

White, Alexander. 2013. "Search Engines: Left Side Quality versus Right Side Profits." *International Journal of Industrial Organization* 31 (6): 690–701.