UCD CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2020

Hops, Skip & a Jump: The Regional Uniqueness of Beer Styles

Bernardo S Buarque, Ronald B Davies,
Ryan M Hynes and Dieter F Kogler,
University College Dublin

WP20/31

December 2020

UCD SCHOOL OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD DUBLIN 4

Hops, Skip & a Jump: The Regional Uniqueness of Beer Styles

Ryan M. Hynes[1], Bernardo S. Buarque, Ronald B. Davies and Dieter F. Kogler[2]

JEL classification: Q10; R11

Keywords: Beer, Economic Geography, Network Analysis

Abstract

Perhaps more than any other product, beer evokes the place it was made. Weißbier and Germany, dubbels and Belgium, and most of all, Guinness and Ireland. Part of what makes these beers so memorable is what sets them apart and gives them their 'taste of place'. Many studies have tried to place that taste, and due to a lack of detailed data, have relied largely on qualitative methods to do so. We introduce a novel data set of regionalized beer recipes, styles, and ingredients collected from a homebrewing website. We then turn to the methods of evolutionary economic geography to create regional ingredient networks for recipes within a style of beer, and identify which ingredients are most important to certain styles. Along with identifying these keystone ingredients, we calculate a style's resiliency or reliance on one particular ingredient. We compare this resiliency within similar styles in different regions and across different styles in the same region to isolate the effects of region on ingredient choice. We find that while almost all beer styles have only a handful of key ingredients, some styles are more resilient than others due to readily available substitute ingredients in their region.

[1]Corresponding author: rhynes57@gmail.com

[2]Hynes and Davies: School of Economics, University College Dublin; Buarque and Kogler: Spatial Dynamics Lab, University College Dublin

We thank participants at the UCD/UCLA Seminars in Economic Geography Series for useful comments. All errors are our own.

1 Introduction

When we think about beer, we think about location; dubbels and Belgium, Weißbier and Bavaria, Guinness and Dublin. In an era of globalization, why do we still have these distinctions? Is there something about the way these beer styles are made that make them 'local' to a region, even though they are produced all over the world? If the answer to this is yes, then geography clearly still matters for brewing.

To investigate the link between styles, ingredients, and places, we borrow economic geography's approach to studying knowledge and place. In that literature, the co-occurrence of technological codes on patents is frequently used to regionalize knowledge networks and link technological combinations to place (Feldman et al., 2015; Rigby, 2015). We use information on beer recipes in a similar way. Just as technologies are combined into a patent, ingredients are combined in a beer recipe. Like regional knowledge networks, we create regional style networks by attributing a style to its traditional place of origin (e.g. American IPA or Kölsch). This mapping allows us to compare close styles across different physical locations (with American versus English IPA our illustrative example) as well as across different styles within the same physical location (with a comparison of Kölsch and Munich Helles as our example).

This exercise reveals several insights. First, the ingredient network is far closer within regions than within styles. In other words, different styles within the same country like German pale ale Kölsch and lager Helles have greater overlap in ingredients than styles within the same family, like the American and English versions of IPA. This suggests that even in a globalized world, local ingredients play a critical role in the distinction of a style of beer, providing some evidence behind the desire to attribute authenticity to a physical location via appellations.

In addition, our analysis measures which ingredients are central to a style's definition, how resilient a style is to losing those ingredients, and how a regional dearth or abundance of ingredients influences a style's development. This analysis provides an empirical basis for discussing concerns such as the susceptibility of some crops to pests or climate change and the knock-on implications for the sustainability of certain beer styles (Yool and Comrie, 2014). We thus contribute to the existing literature on the geography of beer by bringing this detailed data set to bear on a classic question of the importance of place. We also contribute to the literature on Evolutionary Economic Geography (EEG) by applying its methods to product-level data and identifying how intermediate components and their sources affect product resilience.

There is a large body of existing literature discussing the regional aspects of beer. Mittag (2014) considers the development of geographic appellations of beer from their origins as brewery names to their inclusion as distinct styles in the official Beer Judges Certification Program (BJCP) Style Guide. These appellations and styles are an implicit acknowledgement that for its seemingly large variety, a majority of the world's beer styles originate from only a handful of countries. In short, Mittag's assertion is that place is critical to beer. Yool and Comrie (2014) qualitatively consider this concept and the unique combination of ingredients that give beer its 'taste of place', warning that climate change could threaten beer ingredients in sensitive growing regions. Kind and Kaiser (2020) voice similar concerns over Germany's Hallertau region, and note the general sensitivity of hops to extreme weather. Knudson et al. (2020) chronicle the dominance of the Pacific Northwest in US hop production, but also note recently renewed production efforts in several other US regions as demand for more local ingredients increases. We provide rigorous empirical evidence in support of these conclusions. Sewell (2014) provides a historical summary on the spatial diffusion of beer from its origins in the Fertile Crescent, to Ancient Rome, Europe, the United States, and to modern times with the rise of microbreweries.

Microbreweries are an especially well-studied topic in the literature. Microbreweries represent the combination of innovation, entrepreneurship, and rapid growth. Elzinga et al. (2015) chart the growth of the American craft beer industry from 1979-2012, noting that craft breweries tend to appear in geographic clusters. Dennett and Page (2017) reach similar conclusions, and find that two distinct geographic clusters drove the recent expansion in the London craft brewing industry. Wojtyra et al. (2020) find the same clustering of microbreweries in hot spots they identify in Eastern Europe.

Why do these microbreweries gather in one location? Is it access to trade, small business-friendly incentives, or ingredients? Flack (1997) provides one possible answer: neolocalism. The sense of place craft beer evokes is its main differentiator and selling point. Gatrell et al. (2018) go one step further and suggest that craft and microbreweries use "place, practice, and region" to create a strong spatial brand that is appealing to consumers. Cabras and Bamforth (2016) note that consumers often associate local craft breweries with higher quality beer, even when those breweries actually lag behind their larger counterparts in quality control and consistency. Microbreweries clearly rely on place to differentiate their brand and their products. It appears that the same is true for styles as well. Yet for all this research on beer and place one crucial ingredient is missing - the ingredients themselves. This is the gap we fill.

There has yet to be an empirical study of beer ingredients and regional variation. That is likely because it is difficult to get relevant and inclusive data on the subject. Our first contribution is assembling data on beer recipes, styles, ingredients, and their locations. We use nearly 100,000 recipes created by home brewers, craft brewers, and microbreweries. We normalize these recipes and extract ingredients, styles, and their locations to create ingredient networks for each unique style of beer. These networks allow us to quantitatively explore what sets beer styles apart, and to see if there is a central ingredient responsible for a given recipe's 'taste of place'. We do so by borrowing techniques from EEG and constructing style networks composed of all the recipes and ingredients used in beers of a given style.

Our second contribution plots beer in space, and considers which keystone ingredients separate and define different styles. We compare our beer style networks with one another to identify which ingredients are most central to a given style network and therefore define the style. This contribution is a novel application of EEG methods, which typically utilize patent, publication, or skills data as opposed to product-level indicators in generating knowledge networks (Clark et al., 2003; Kogler, 2016). Expanding these methods is crucial to advance the field as a whole because knowledge exists not just in the ivory tower, but also in the everyday products all around us. Even something as deceptively simple as beer is full of complex relationships and is ripe for detailed analysis.

We take inspiration from the knowledge space methodology of Kogler et al. (2013, 2017) and Buarque et al. (2020) which maps patent technology codes in space and create regional knowledge networks. We instead create style networks, where individual nodes are ingredients used in a style and edges are any two ingredients' co-occurrence with one another within recipes. Styles themselves have strong historical ties to specific regions, and are often named after and incorporate ingredients from the specific region where the style was first produced. The historical origins of styles provide our link between styles and their recipes and geographic regions. We are the first to apply this analysis to recipe data, and hope that this novel application inspires others to do likewise.

We are able to create highly detailed recipe-style networks because we collect the weights and measures of individual ingredients within a recipe. These allow us to properly weight edges between ingredients based on their relative proportions within the recipe.[1] After preparing these data, the network algorithm minimizes total network path length, placing the most important and frequently used nodes at the center of the network. We can then easily measure which of the ingredients is the most important to a region. We can also observe how resilient a given style is to the loss of key ingredients, something forewarned by Yool and Comrie (2014) as well as Kind and Kaiser (2020). We measure resilience by deleting key ingredients from the network and observing changes in overall network characteristics.

While this network analysis tells us which ingredients are most important to a network and which ingredients set regional styles apart, we are also interested in how geography shapes a style's resilience and reliance on ingredients. We turn to the con-

[1]Note that most EEG analyses are unable to identify the relative importance of an ingredient in the development of a novel product or process, for example, the varying importance of individual technology codes listed on a single patent document. We are able to do so in the present study because we capture the weight and volume of ingredients.

cepts of relatedness and unrelated variety to measure both breadth of ingredients a style uses, as well as the depth of those ingredients' potential substitutes (Whittle and Kogler, 2020). We find that New World styles generally make use of a larger variety of easily substituted ingredients and are much more resilient than Classic styles because of this.

In summary, we bring highly-detailed micro-data to longstanding questions in the geography of beer literature. We also marry this literature with analysis from EEG and innovation studies. In doing so we shed light on old questions and pave the way for others to ask and answer new ones.

The rest of this paper proceeds as follows: Section 2 discusses how we fetch, parse, and normalize the recipe-level data. Section 3 transforms recipe ingredient data into style networks. Section 4 introduces eigenvector centrality, our main measure of ingredient importance. Section 5 details our targeted deletion strategy. Section 6 defines the ability of certain styles to weather losses of key ingredients. Section 7 posits that geography and the abundance of ingredients is a key determinant of resiliency. Section 8 concludes.

2 Data Collection and Mapping

We gather data on 126,256 beer recipes and map them to individual styles, which in turn can be historically linked to countries, regions, and even cities. We use the authoritative BJCP Style Guide to define broad styles of beer, then match beer recipes to styles. We get our beer recipes and their component ingredients by downloading BeerXML files from BrewersFriend.com. BrewersFriend allows home brewers and small craft breweries to record and manage their recipes. Recipe ingredients are broken down into hops and malts, each of which detail the types and amounts of ingredients added to the recipe. Figure A (Appendix) provides an abridged example of the BeerXML file for one such recipe.

BrewersFriend allows recipes to be made publicly accessible or otherwise marked private. The 126,256 public recipes on BrewersFriend form the basis of our sample. We download these public recipes in BeerXML format, then parse the ingredients in each of the five categories into separate tables. Once parsed, we spend considerable effort disambiguating ingredient names so that they may be matched to multiple recipes.[2] We then turn to refining our sample.

We first restrict our sample to include only recipes using whole ingredients. Some recipes in BrewersFriend use pre-mixes from brewing kits that already combine ingredients and therefore offer little information about the choice or combination of ingredients. This restriction leaves us with 109,015 unique recipes, or 86% of our original sample. We then turn to regionalizing our recipes through their styles.

Each recipe is associated with a single official BJCP style. BJCP styles are an international standard used to group and evaluate beers at brewing competitions worldwide. Most BJCP styles are associated with a given country and region, for example Kölsch is a specific BJCP style originating from Köln in Germany. We group our recipes into 144 different BJCP styles, and drop 3,159 recipes that do not specify a style. We drop these BJCP "Specialty Beers" styles including mead, cider, and other non-beers and lose an additional 4,821 recipes (4% of our remaining total). We are left with 101,034 recipes covering 111 styles. Table B (Appendix) lists these styles and the number of recipes in each.

The distribution of recipes across styles is highly skewed. Two styles, American IPA and American Pale Ale, represent more than 25% of all recipes. This may represent an underlying bias in our data as BrewersFriend is based in the United States, or it could also reflect the tremendous popularity of these styles.[3] However, there are thousands of international users of BrewersFriends and over 100 styles with at least one thousand recipes each. Figure 1 shows the distribution of recipes by individual style.

There are a handful of styles, such as New Zealand IPA, that only have one recipe

[2]For example, one recipe may use "CaraPils" malt and another "carapils" malt, even though these are the same underlying ingredient.

[3]As a robustness check, we randomly draw a sub-sample of American IPA recipes in proportion to the number of recipes in the styles we compare with American IPA. Our results are largely unchanged, so we present the full network in our comparisons below.

Figure 1: Count of Recipes per Style

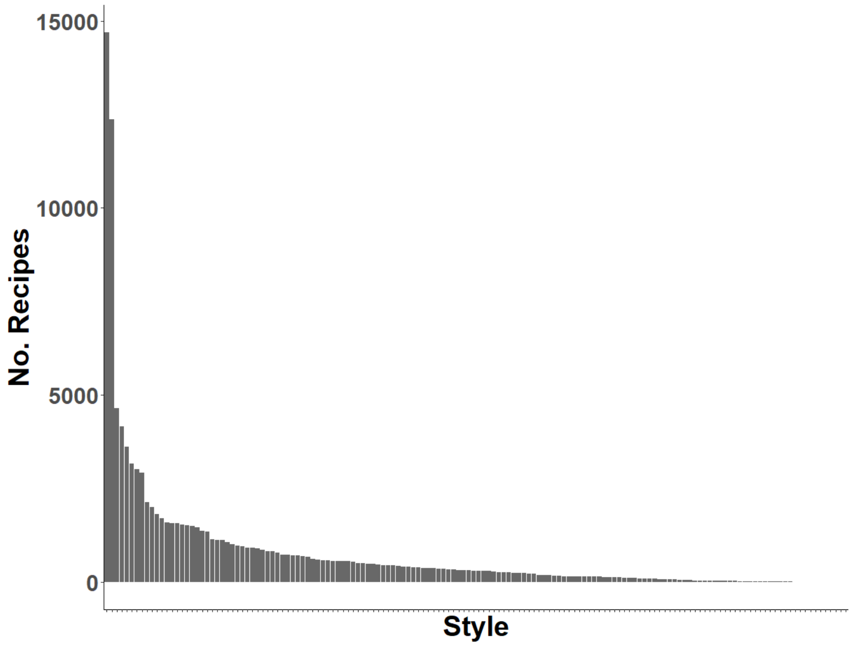

associated with them. To ensure adequate variation within styles, we further restrict our sample to styles that have at least 100 recipes. We lose only 596 recipes (less than 1% of our sample) with this additional restriction, but do miss out on a few valuable regional styles like Rauchbier, which is particular to the town of Bamberg in Germany. After cleaning and regionalizing our style sample, we have 100,438 recipes covering 90 styles and spanning 13 countries. We map these styles into 25 regions with varying levels of precision[4]. We then turn to our ingredient types of interest: hops and malts[5].

We identify 4,882 different malt names across our sample, however not all of these malts are truly unique due to minor variations in their names. We disambiguate these malts by first removing all nationality and company information from the name[6]. We then remove special characters and lowercase all names. We fuzzy match our cleaned

[4]Some beers like the California Common can be located to a specific city and even a particular brewery: Anchor Brewing in San Francisco. Others have less precise origins. The American IPA is primarily attributed to the West Coast of the United States, but is also fairly ubiquitous across the country. Finally, most British beers can only be mapped to the national level, i.e. Scotland or England.

[5]BrewersFriend.com provides five categories of ingredients: Hops, Malts, Yeasts, Waters, and Miscellaneous. We focus on hops and malts because they: 1) are arguably the most important ingredients in recipes; 2) are almost always combined with different varieties in recipes, as opposed to yeasts; 3) are the most readily identifiable and easy to accurately localize.

[6]For example, "US - Castle Malting - Pilsner Malt" simply becomes "pilsner".

Table 1: Recipes and Styles by Country

Country	Recipes	Styles	Avg Color	Avg ABV	Avg IBU	Avg OG	Avg FG	Avg Carb.	Avg Malts	Avg Hops	Avg Oth.
United States	56602	36	11.12	6.10	50.15	1.06	1.01	0.94	2.71	2.52	1.38
United Kingdom	15770	34	24.16	6.05	37.30	1.06	1.02	0.80	3.35	1.88	1.47
Belgium	11577	17	8.95	6.43	25.81	1.06	1.01	1.01	2.51	1.85	1.42
Germany	8307	34	10.07	5.66	21.64	1.06	1.01	1.15	2.58	1.56	1.30
Specialty Beer	4164	37	12.37	6.51	27.39	1.06	1.01	0.61	2.90	1.76	1.58
Ireland	2778	3	18.60	5.20	28.00	1.05	1.01	0.95	3.27	1.68	1.17
Russia	1302	1	47.48	9.90	60.98	1.10	1.02	0.60	4.30	2.01	1.55
Czech Republic	857	5	6.23	5.52	35.37	1.05	1.01	0.95	2.24	1.68	1.08
International	525	3	7.24	5.62	26.62	1.05	1.01	1.26	2.21	1.87	1.22
Austria	492	1	11.08	5.33	25.20	1.05	1.01	0.89	2.85	1.79	1.10

list of malts back to the recipes and confirm the matches by hand. Like styles, we remove infrequently used malts appearing in fifteen or fewer recipes. We are then left with 170 unique malts used in 99,943 recipes. Table C (Appendix) lists all these disambiguated malts.

Like Malts, we begin with a list of 5,023 Hops that appear at least once across our recipes. We repeat the name normalization process above, removing brand names and indications of origin. We once more confirm these results by hand, paying particular attention to code names and translations. For example, one common hop, Saaz, is known in the Czech Republic as Žatec, the name of the town where it is produced. After normalizing we perform a fuzzy match to a list of well known hops provided by both Barth-Haas and Hoplist.com. Barth-Hass is one of largest producers of hops worldwide and has developed a ubiquitous Tasting Guide detailing the flavor profile, alpha acid, and location of global hops (BarthHaas, 2018). Hopslist.com similarly maintains a global reference of hops and their locations (Healey, 2016).

We disambiguate our 5,023 hops from our recipes to just 229 global hops from the Barth-Haas and Hopslist.com lists. Unmatched hops are almost all due to misclassifications such as listing fruit or spices as hops, or other user data-entry errors when creating the recipe. We similarly restrict our sample to hops appearing in 15 or more recipes to ensure adequate variation across our sample. We lose only 362 recipes with this restriction, leaving us with 229 unique hops used in 92,813 different recipes. Table D (Appendix) lists these disambiguated hops.

After parsing, disambiguating, and cleaning our sample we are left with 92,813 recipes made from 170 malts and 161 hops across 90 different styles. Table 1 summarizes these data. We now use these data to create recipe-ingredient networks for each style.

3 Beer Style Networks

We create an ingredient co-occurrence network for all 90 beer styles in our sample. Each recipe represents a unique combination of hops and malts, at the extensive margin if a particular ingredient appears in a recipe, and at the intensive margin based on the relative proportions used of each input. These style networks are graphical representations of the distinct combinations of ingredients and their volumes.

Each ingredient is a node in the style network. We draw an edge between two ingredients whenever they co-occur in the same recipe. Each edge is weighted in proportion to the amounts used in the recipe. For example, if a recipe uses 1kg of Pale 2-Row malt for every 100g of Chocolate malt, we value the edge between these two ingredients as 1/10.[7] Because every beer belongs exclusively to one style, we can combine the nodes and add their weighted edges to form unique style networks. If the same ingredient pair appears in more than one recipe of the same style, we sum up their weights.

Our style networks describe the relationship between the ingredients used in every beer recipe of a given style. The networks allow us to visualize the unique combinations of ingredients that make up a beer style. We can also represent these style-ingredients relationships algebraically:

$$S_{ij} = \begin{bmatrix} s_{11} & s_{12} & \dots \\ \vdots & \ddots & \\ s_{n1} & & s_{nn} \end{bmatrix} \tag{1}$$

where S_{ij} is the style's adjacency matrix and every entry s_{ij} measures how often ingredients i and j appear together in recipes of the style, weighted by their relative proportions. The adjacency matrix above can also be visualized as a style network.[8]

Figure 2 plots two such style networks. Panel A shows the style network for Amer-

[7] One does not usually observe the volume of each input used in the end product when using patents or other data sources to build co-occurrence networks. Instead, this literature typically weights edges based on the shares of the node. For example, if four technological codes appear in the same patent each gets a weight of 1/4. For this reason, we also reproduce our analysis weighting the edges of the style networks by the ingredients' shares. Our results are robust to using this more common weighting method.

[8] The adjacency matrix, edge list, and networks are different ways to represent the same relationship between nodes and edges. We provide definitions for all three in the Appendix. See primary references Wasserman et al. (1994) and Barabási et al. (2016) for further information.

ican IPA, the most popular style in our sample. Panel B shows the style network for Kölsch, a beer style named for the Germany city where it was first created, which perhaps best captures the regional nature of styles.[9] We create both graphs using the Kamada-Kawai force-directed drawing algorithm, which minimizes total path length and places ingredients that are commonly used together next to one another (Kamada et al., 1989). Likewise, Kamada-Kawai puts the most connected nodes at the center of the network. The size of each node is proportional to that node's degree centrality, or how many connected links a node has. The width of each edge is proportional to the weights of the ingredients' as they co-occur in recipes. Hops are colored in green and shaded by their alpha acid intensity, a proxy for bitterness. The darker the green, the more bitter the hops. Malts are colored brown and shaded by their European Beer Convention (EBC) coloration.[10]

These two styles and their graphs are quite different. American IPA uses many more unique ingredients than Kölsch (321 nodes against 175). American IPA's ingredients are also more connected to one another with more than 17,000 total edges between its nodes, each of which has 108 edges on average. Kölsch, on the other hand, has only 2,000 total edges and an average of 20 edges per node. The American IPA has a very high network density, which is the number of actual edges between nodes out of all theoretically possible edges. In fact, the American IPA has a relatively high network density of 0.34 or 34% of all possible edges, while the Kölsch has a network density of only 0.13.

The American IPA network seems to be more robust and complex than that of Kölsch. The American IPA network includes more ingredients with stronger connections between them. Still, one might argue that because our sample of American IPA recipes is much larger than any other style, and more than ten times greater than Kölsch, we misrepresent its network connectivity.[11] Nevertheless, these differences ex-

[9]Since 1997, Kölsch holds a Protected Geographical Indication (PGI) within the European Union.

[10]EBC coloration is a grading scale based on the color a particular malt imparts on a beer. Pilsners and other light beers have an EBC of 4, whereas darker malt beers such as stouts have an EBC of 70.

[11]We test if the differences between the American IPA and Kölsch are the result of sample size. To do so, we take 1,000 random sub-samples of American IPA consisting of 1,000 recipes each, approximately the

Figure 2: Example Networks

(A) American IPA

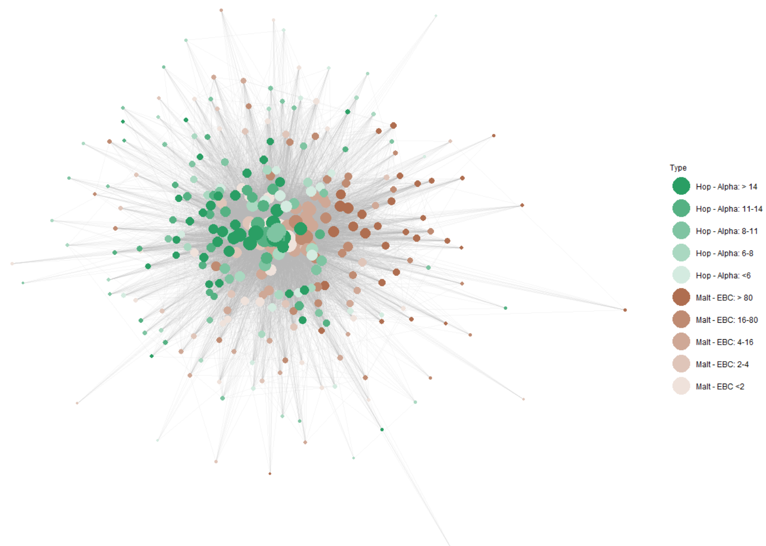

Nodes	**Density**	**Diameter**	**Clustering**	**Avg. Degree**
321	0.34	1.20	0.54	108

(B) Kölsch

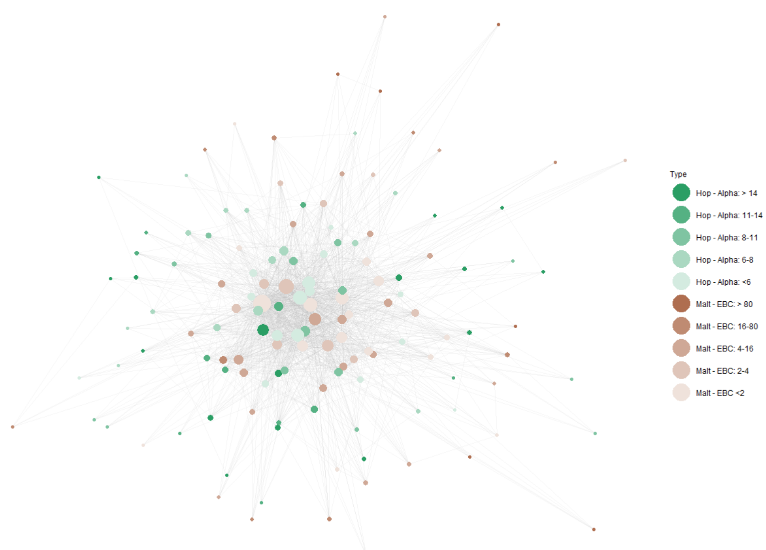

Nodes	**Density**	**Diameter**	**Clustering**	**Avg. Degree**
175	0.13	1.50	0.39	23

same number of Kölsch recipes. Although on average the American IPA random sub-sample networks are not as connected as the full sample American IPA network, they remain more connected than the Kölsch network. The random sub-samples have more nodes (22) and edges between them (4,600), more than the Kölsch network. The sub-sample networks also have a higher density (0.17), average clustering coefficient (0.40), and degree (40), as well as a shorter diameter (1.4).

ist across all styles in our sample, and are evident even when considering networks with a similar number of recipes.

For example, California Common, another typical beer from the American West Coast, has a similar number of recipes as Kölsch at about 900 each. Despite the similar number of recipes, California Common lists more ingredients (220 nodes) and a higher average number of edges per node (38). California Common's average clustering coefficient is 0.46 compared to Kölsch's 0.39, meaning California Common includes relatively more "three-way" connections between ingredients. The California Common's network is also smaller than Kölsch's in the sense that it takes fewer steps to traverse the network. Indeed, California Common's maximum shortest path, or diameter, is 1.0 compared to Kölsch's 1.5.

These style networks provide us with a tractable method to visualize and model the relationship between recipes and ingredients. We now use these models and their properties to identify key ingredients, resilience, and relatedness across beer styles.

4 Eigenvector Centrality

While there is clearly significant variation among styles and their networks within our sample, our main goal is to identify which key ingredients set these networks apart and give beer styles their unique tastes. In other words, we are looking for the most important ingredient nodes in a given style network. We turn to eigenvector centrality as a measure of each node's relative importance within a network. We follow the seminal work of Bonacich (1972) and calculate eigenvector centrality as the weighted sum of the centrality of all adjacent nodes. Mathematically we can express eigenvector centrality as:

$$\lambda c(v_i) = \sum_{j=1}^{n} s_{ij} c(v_j) \tag{2}$$

where λ is the eigenvalue scale factor, $c(v_i)$ represents the centrality score of node vector v_i and s_{ij} is the weighted edge between nodes i and j. Algebraically this represents every element in the adjacency matrix (S_{ij}) associated with our style networks.

Eigenvector centrality differs from traditional degree measures of importance because it also accounts for the relevance of a node's immediate neighbors. As Ruhnau (2000) explains: "centrality of nodes does not only depend on the number of its adjacent nodes but also their value of centrality" (p.360). Eigenvector centrality awards points for being linked to very central nodes even if the node itself has just a few connections. For this reason, it is often used in social sciences to measure the influence of agents (Abbasi et al., 2011; Li et al., 2016; Parand et al., 2016).

Table 2 shows the top ten ingredient nodes by eigenvector centrality in our original American IPA and Kölsch networks. We normalize the centrality scores between zero and one, such that the most central node in each network will always have a score of one.[12]

Table 2: Top Ten Nodes by Eigenvector Centrality

American IPA		
Ingredient	Type	Eigenvector
Citra	Hop	1.00
Pale 2-Row	Malt	0.91
Cascade	Hop	0.77
Amarillo	Hop	0.75
Simcoe	Hop	0.74
Centennial	Hop	0.74
Mosaic	Hop	0.63
Columbus	Hop	0.49
Chinook	Hop	0.47
Maris Otter	Malt	0.36

Kölsch		
Ingredient	Type	Eigenvector
Pilsner	Malt	1.00
Hallertau	Hop	0.73
Tettnanger	Hop	0.54
Saaz	Hop	0.44
Vienna	Malt	0.39
Hersbrucker	Hop	0.36
Wheat	Malt	0.29
Perle	Hop	0.27
Pale 2-Row	Malt	0.22
Magnum	Hop	0.20

Once again, there are considerable differences between these two styles, this time in key ingredients. The most central nodes for American IPA are mostly bittering hops with high-intensity alpha acids from the Yakima Valley in Washington State: Citra, Cascade, Amarillo, Centennial, etc. On the other hand, Kölsch relies heavily on aromatic hops traditionally found in Pilsners and Lagers from the Bavaria and Bohemia regions such as: Hallertau, Tettnanger, and Saaz. There are likewise significant dif-

[12] We consider hops and malts together as both are fundamental ingredients to recipes, which use each in different combinations.

ferences in eigenvector centrality between the top ten ingredients of both styles. The distance between the first ranking ingredient in Kölsch and the rest is much greater than that in American IPA, implying the German style relies more heavily on a single malt source: Pilsner. Figure 3 further illustrates this difference and plots the histogram of eigenvector centrality for all ingredients in both beer styles.

Figure 3: Eigenvector Centrality Distribution within Styles

(A) American IPA

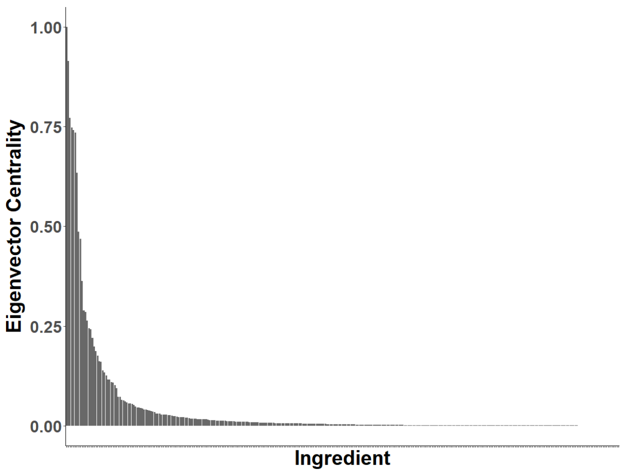

(B) Kölsch

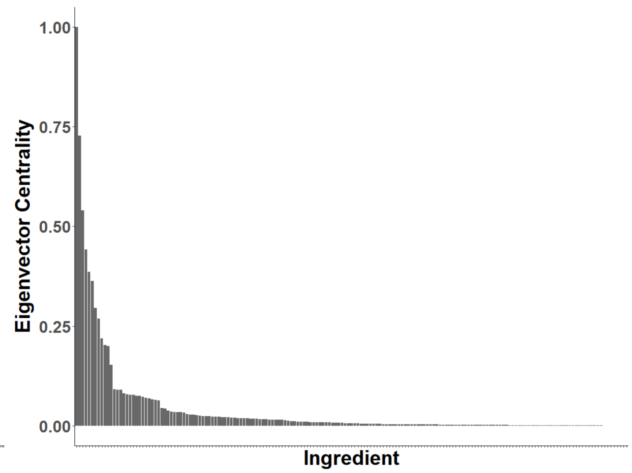

Note: These plots are the eigenvector centrality distributions for every ingredient in the American IPA and Kölsch networks. The x-axis lists ingredients ranked by centrality score. The y-axis is the eigenvector centrality score. We measure centrality according to the eigenvector formula developed by Bonacich (1972). We normalize centrality scores between one and zero, such that the most central node always has a centrality score of one.

Both histograms in Figure 3 show signs of long tails common in power-law and Pareto distributions, which confirm that our beer networks display scale-free properties prevalent in many social, biological, and physical systems (Newman, 2005). In scale-free networks, there are often a small number of highly connected nodes with most other nodes having little to no edges. This unequal distribution persists even when the system expands or contracts, hence the name scale-free.

Because the number of edges per node is so skewed, a common trait across scale-free networks are their resiliency to "errors" or the loss of nodes. Because most nodes have few connections, deleting a random node from a scale-free network does little to change the network's overall structure and function. Conversely, scale-free networks are extremely vulnerable to "the selection and removal of a few nodes that play a vital role in maintaining the network's connectivity" (Albert et al., 2000, p.379).

The concepts of error tolerance and attack vulnerability are fundamental for designing and understanding communication networks such as the World Wide Web. Beer is definitely not the Internet, so instead it helps to imagine a scenario where due to climate change or diseases we are no longer able to produce one or two varieties of hops. Depending on the centrality of these lost hops in the style network, we ought to expect different effects on the network structure and number of feasible recipes. If the lost hops are very central to the style network, we would expect its structure to change significantly. If instead the hop is peripheral, the network structure and its observable characteristics would not change much at all. To put this idea idea into practice, imagine the world is no longer able to produce Citra hops. Kölsch beers would not fundamentally change, whereas the network structure and frontier of possible recipes within the American IPA network would be significantly reduced. We explore this network resiliency and sensitivity to particular ingredients in Section 5 below.

5 Stress Test

Rather than just observing a given node's centrality in a network, we can ask: what if that node had never existed in the first place? This approach is referred to as network fragility or resiliency analysis and allows us to measure aggregate network statistics like density, path length, and centrality as a function of one particular node. We follow Albert et al. (2000) and Toth et al. (2020) and iteratively remove nodes from our style networks and recalculate key network statistics to measure how a network changes in the absence of a given node. In our case, this approach reveals how sensitive a given beer style is to losing any one ingredient, which in turn reveals that ingredient's importance to the style.

We run this stress test in two ways. First, we delete nodes in rank order according to their eigenvector centrality. Second, we delete nodes at random as a baseline comparison. To truly randomize this deletion process, we run 10,000 iterations of random deletions for each network and report the average changes in network statistics. We provide a glossary of these network statistics and their definitions in Table A (Appendix).

Figure 4 shows the consequences of both targeted and random deletion in the American IPA network. Panel 4A shows the resulting network after targeted deletion of 40%, 60%, and 80% of the most central nodes according to eigenvector centrality. Panel 4B shows the same 40%, 60%, and 80% deletion, this time removing nodes randomly. Like Figure 2, we use the Kamada et al. (1989) network plotting algorithm which places the most central nodes in the middle. To better visualize the effects of deleting nodes, we fix the network at its original layout, then remove nodes and edges from it. However, we properly re-scale the network after each deletion when re-calculating network statistics. As before, the node sizes are proportional to the weighted number of connections, and their colors depend on the ingredient type and intensity.

Panel 4A clearly shows the sensitivity of the American IPA network to targeted deletion. In contrast, Panel 4B shows American IPA's relative resilience to random deletion. Even if we delete 40%, 60% or 80% of the nodes, the resulting networks from the random attacks have more connections and shorter paths relative to the targeted attack networks. To measure how much variation we obtain from the deletions, we compute four key network statistics and compare them to the full network. We reproduce the randomization order 10,000 times and save the density, diameter, average clustering coefficient, and average degree from the resulting networks. Figure 5 shows the distribution of the absolute percentage change our four network statistics after randomly deleting 50 nodes. We also highlight the changes in those statistics from a targeted deletion of 50 nodes with a dashed red line. Figures B and C (Appendix) repeat this targeted and random deletion exercise for Kölsch to much the same effect. The effects of the targeted attack are clearly much greater than its random counterpart. Even though the American IPA is the largest style in our sample and perhaps the

Figure 4: Stress Test - American IPA

(A) Targeted Attacks

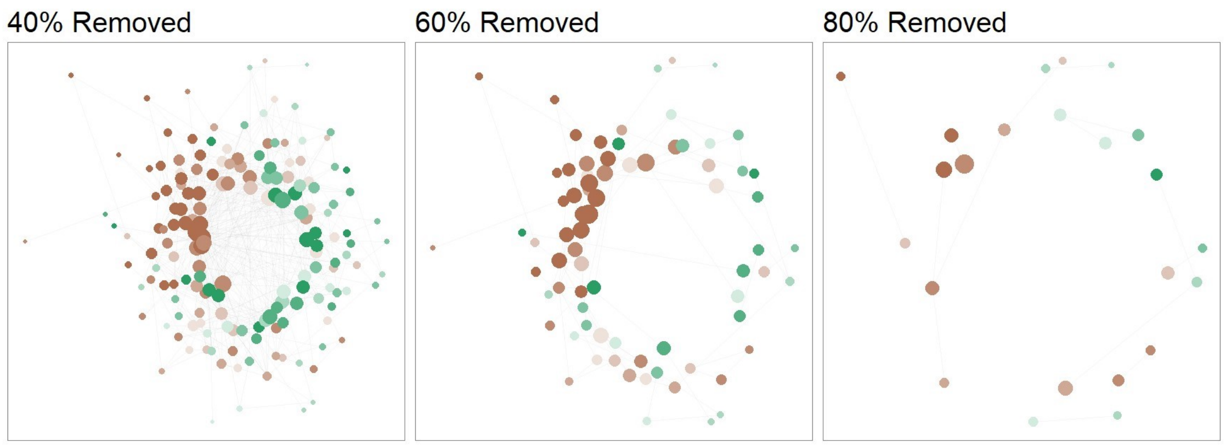

(B) Random Attacks

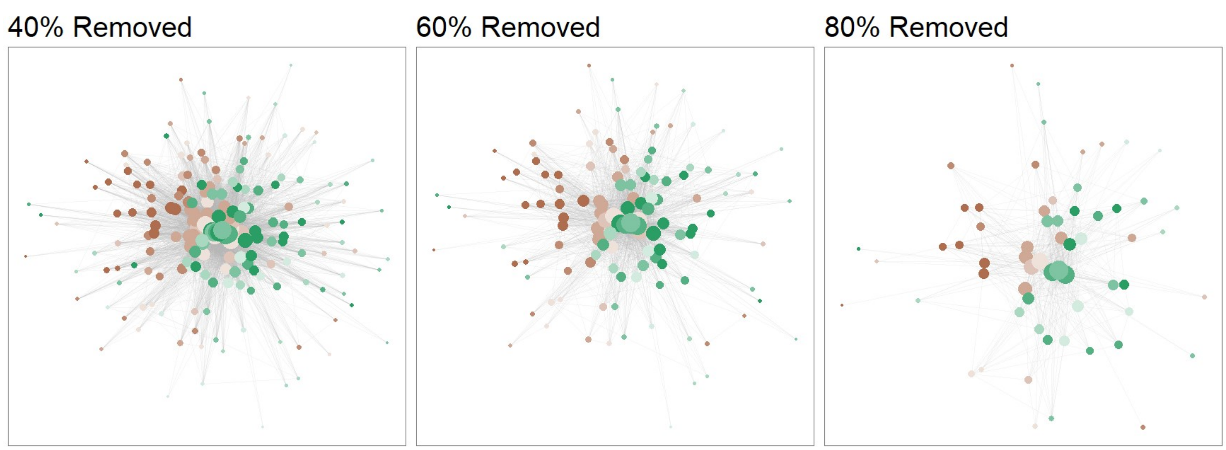

Note: Both panels depict the impact of removing 40%, 60% or 80% of the nodes from the American IPA network. Panel 4A shows the effect of targeted deletion according to eigenvector centrality. Panel 4B shows a random attack where we nodes are deleted in random order.

most connected network out all styles, it relies on just handful of keystone ingredients without which the entire style network crumbles. These keystone ingredients are what differentiate styles and create a unique, identifiable flavor.

Figure 5: Resiliency Against Random Removal

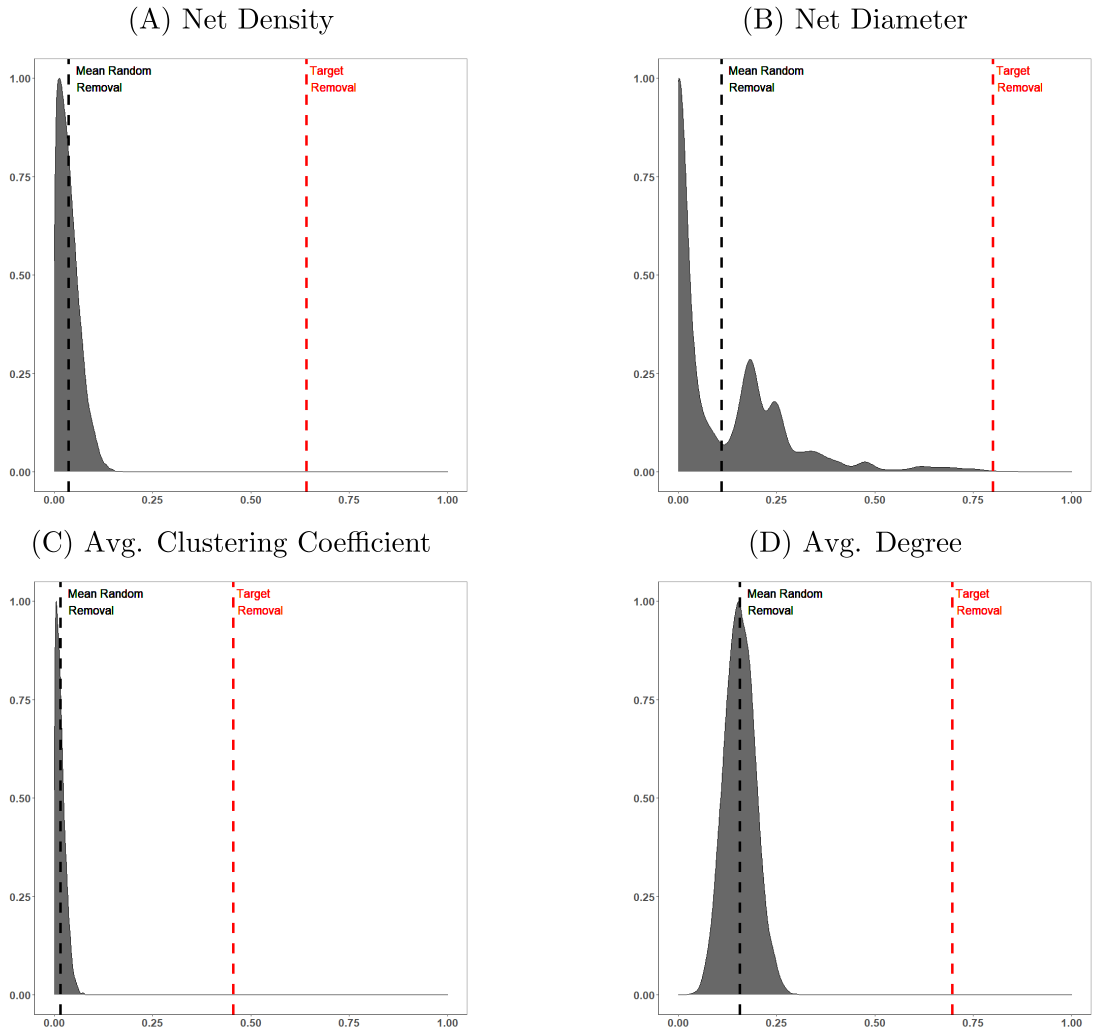

Note: These figures plot the probability density distribution of the effect of deleting 50 nodes from the American IPA network. The y-axis is the probability density scaled between 0 and 1, such that the most frequent effect is equal to 1. The x-axis is the absolute value of the percentage change of a given network statistic. Density refers to the number of edges out of total possible links. Net Diameter is the maximum shortest path. Clustering coefficient is the fraction of total three-way connections out all possible ones. Average degree is the average number of edges each node has. A network is no longer connected when Average Degree falls below one. The dashed black line is the average effect of 10,000 random deletions. The dashed red is the effect of targeted deletion according to eigenvector centrality.

6 Resiliency

A common feature across all beer styles is their high dependence on a few key, central ingredients. All style networks show scale-free properties and thus are vulnerable to the failure of just a few ingredient nodes. However, there is significant variation in ingredient dependence across styles. Beer styles are not equally resilient and deleting the most central nodes in one style might have a more powerful effect than in another.

Let us return to our original example and compare the network structures and eigenvector centrality distributions of American IPA and Kölsch. American IPA is more resilient because it has a larger number of connections and many ingredients with a relatively high eigenvector centrality. As such, it can afford to lose more critical nodes than Kölsch. Is this a unique attribute of American IPA alone, or common to the larger family of IPA styles?[13] To find out, we compare American IPA within and across style families. Table 3 introduces two new styles, English IPA and Munich Helles, and shows the network consequences of targeted deletion of the top fifty ingredients for all four styles. Figure 9B (Appendix) plots the ingredient networks for these two additional styles.

All four example networks experience a loss in connectivity after deleting the top five, ten, twenty, or fifty most central nodes. After deleting the top 20 nodes, every network is nearly half its original size by density or average degree. Likewise, network diameter nearly doubles after removing the top 20 nodes, meaning all four networks are becoming less connected and more difficult to traverse. Despite these similar trends, Table 3 also shows variation within each style's ability to withstand shocks. American IPA experiences the largest overall drop in average degree after deleting fifty nodes, yet remains more connected than the full Kölsch or Munich Helles networks. Moreover, this is not just a function of the IPA style family, as closely-related English IPA does not exhibit the same resiliency. What makes the American IPA so much more robust than other networks? It is clearly not only a function of sample size, but rather is the result of the style's relative fungibility of key ingredients. American IPA has greater availability of close substitutes because it makes use of more diverse ingredients.

Turning back to our hypothetical where Citra hops go extinct, American IPA still

[13]The BJCP also defines several 'Style Families' that group multiple related styles. These families are listed in Table B (Appendix).

Table 3: Network Resiliency

Style	Nodes Deleted	Density	Diameter	Avg. Clust. Coeff.	Avg. Degree
American IPA	0	0.34	1.20	0.54	108.50
	1	0.33	1.20	0.54	106.04
	5	0.31	1.50	0.53	96.60
	10	0.27	1.50	0.52	85.74
	20	0.22	1.55	0.50	66.80
	50	0.12	2.16	0.37	32.84
English IPA	0	0.19	1.16	0.45	49.17
	1	0.18	1.16	0.44	47.59
	5	0.16	1.70	0.43	41.77
	10	0.14	1.47	0.40	35.38
	20	0.10	1.91	0.35	25.10
	50	0.05	2.45	0.23	10.64
Kölsch	0	0.13	1.50	0.39	22.96
	1	0.12	1.81	0.40	20.89
	5	0.09	2.00	0.38	16.00
	10	0.07	2.33	0.34	11.71
	20	0.05	2.42	0.27	7.62
	50	0.01	3.30	0.30	1.85
Munich Helles	0	0.16	1.85	0.42	15.39
	1	0.14	2.15	0.43	13.63
	5	0.10	2.14	0.37	9.67
	10	0.07	2.15	0.35	6.68
	20	0.04	1.97	0.32	3.23
	50	0.01	0.83	1.00	0.62

has many alternatives with similar traits to choose from. This why the geography of a style is so important. The United States produces more than 60 different types of hops, many of which are very similar to Citra because they are grown in the same regions. In fact, brewers refer to Citra and its sister hops as the '7Cs', which also include: Cascade, Centennial, Chinook, Cluster, Columbus and Crystal. All of the 7Cs are known for their intensity and bright citric flavour. So while the Citra hop is a key ingredient of American IPAs, it is also easily replaceable. It is then important to understand the correlation between a style's resiliency, the availability of related ingredients and the overall diversity of inputs used. We introduce three new variables to measure these factors.

We measure the resiliency of each style network according to Toth et al. (2020), who study the co-occurrence of patent classes and define technological resiliency as the "amount of node removal that a region's technology network could withstand without being fragmented into many unconnected components" (p.13). We use the Molloy and Reed (1995) criterion as the threshold below which a network fragments into many separate pieces. Mathematically the criterion is:

$$\Omega_s = \frac{\sum_{i=1}^{N} k_{is}^2}{\sum_{i=1}^{N} k_{is}} \tag{3}$$

where Ω_s is the resiliency score, or the percentage of nodes removed before the Molloy-Reed criterion falls below two, and k_{is} is the average degree, or number of edges each node in the network has. Having defined a measure of network resiliency, we now introduce two of its key determinants: related and unrelated variety.

EEG discusses the differences between related and unrelated variety and how these properties shape the ability of firms and regions to diversify, innovate, and grow (Content and Frenken, 2016; Boschma, 2017; Miguelez and Moreno, 2018; Rocchetta and Mina, 2019). We borrow these concepts to understand how the availability of substitutes for key ingredients shapes the resiliency of our style networks. We take related variety to represent the presence of similar substitutes e.g. Citra or Chinook, while unrelated variety is a style's ability to source from multiple and distinct products, e.g. Pale 2-Row and roasted barley.

We measure unrelated variety according to Frenken et al. (2007) and we apply the Shannon Entropy formula (Shannon, 1948) to the incidence of ingredients in a style as follows:

$$UV_s = \sum_{i=1}^{N} P_{is} \log_2 \left(\frac{1}{P_{is}} \right) \tag{4}$$

where P_{is} is the probability of finding ingredient i in beer style s. The Shannon Entropy formula applied to our beer styles captures the level of "uncertainty" or "surprise" across each style's recipes. In our style networks, Shannon Entropy measures the likelihood a recipe includes an unexpected ingredient not commonly found in other beers that belong to the same style, as well as how styles source distinct ingredients. For an example of a surprising ingredient, think of using a Chocolate type malt, typically found in dark and robust Stouts, to make an American IPA. Thankfully this unsavory combination is not very common, though it is certainly possible and would contribute towards a larger entropy or unrelated variety for the style.[14]

Frenken et al. (2007) exploit the unique hierarchical structure of employment to distinguish between related and unrelated variety. However, we cannot apply the same approach to our beer recipes as we cannot separate ingredients into hierarchical structures. Instead, we follow Kogler et al. (2013, 2017) and calculate average relatedness of individual ingredients as a measure of related variety.

We first create a global co-occurrence network covering all beer recipes in our sample regardless of style. The global network follows the same structure as the individual styles described in Section 3. We use this network to measure the similarity or relatedness between each ingredient pair. The more often two ingredients appear together across recipes, the more similar they are and the closer their "cognitive" proximity (Nooteboom, 2000). We measure relatedness by standardizing the elements of the adjacency matrix by the square root of the product of the number of recipes in the row and column ingredients of each element:

$$R_{ij} = \frac{s_{ij}}{\sqrt{N_i * N_j}} \tag{5}$$

where R_{ij} measures the relatedness of each ingredient pair, s_{ij} are the elements of the adjacency matrix and measure how often these two ingredients co-occur (weighted by their proportions), and N_i, N_j are the count of total recipes containing each ingredient. Considering the incidence of ingredients within each style to the sum of their proportions we estimate the style's average relatedness as:

$$AR_s = \frac{\sum_i \sum_j R_{ij}(N_i N_j) + \sum_i 2N_i}{P_s(P_s - 1)} \tag{6}$$

[14]Stone Brewing has one such Valentine's-themed example, though Stone gets no love from the Authors for it: https://www.stonebrewing.com/beer/stone-enjoy-ipa-series/stone-enjoy-021417-chocolate-coffee-ipa

where P_s is total count of recipes within each style. Therefore, while unrelated variety measures how much each style sources from various ingredients, average relatedness measures the similarity of ingredients used within a style, where we first estimate relatedness using the global sample of recipes. In other words, average relatedness measures the availability of substitutes for every core ingredient used in a given style. For example, two similar hops like Citra and Mosaic have a relatively high average relatedness of 3.69, whereas two distant hops such as Citra and Hallertau have an average relatedness of just 0.21. The same is true of malts as well. The delicious Pale 2-Row and Chocolate example above also has a mercifully low average relatedness of only 0.21. We conclude that if a style uses more similar ingredients, it will have a higher average relatedness and more readily available substitutes.

These EEG metrics allow us to measure the diversity of ingredients within a style, as well as the importance of having substitutes. It is important to note, however, that these variables are not mutually exclusive. A style could have both high levels of average relatedness and unrelated variety. That is, a style could simultaneously use many ingredients, each with ample substitutes.

After introducing these measures of resiliency, unrelated variety, and average relatedness, we can observe the interplay between them within recipes of a given style. Figure 6 plots this relationship. Styles with higher levels of both related and unrelated variety tend to be more resilient. Taking geography into account, American Styles are more robust than the English, Belgian, or German ones, precisely because of their diverse range of ingredients and easily available substitutes.

7 Geography Matters?

So far we have shown that beer styles are highly dependent on a few central keystone ingredients. It remains to be shown that these keystone ingredients differ across styles, otherwise all beers would rely on the same few ingredients. We now turn to demonstrating how central ingredients vary across styles, and that each beer depends on a

Figure 6: Resiliency vs. Unrelated Variety and Related Variety

(A) Unrelated variety (B) Related Variety

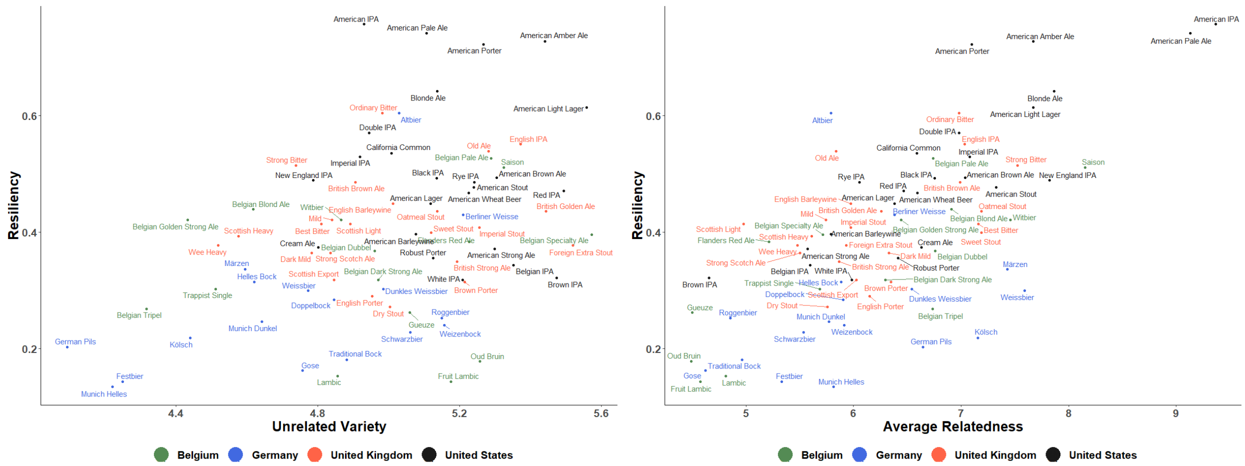

Note: These plots show the correlation between resiliency and either unrelated variety or related variety. The y-axis is resiliency, which we measure as the percentage of nodes a network can lose before fragmenting into many unconnected components. The x-axis is either unrelated variety or average relatedness. We calculate unrelated variety using the Shannon Entropy formula following (Frenken et al., 2007). We calculate related variety following (Kogler et al., 2013). Points are colored according to the country of origin of the beer style.

unique combination of core ingredients. It is these unique combinations that contribute most to a style's network and to its distinctive flavor.

Table 2 highlights that our two sample styles, American IPA and Kölsch, rely on different ingredients with different eingevector centrality scores. Turning to other style networks, we note how distinct nodes are both highly central to the network and also specific to that style. For example, dark roasted barley is the most central component of Irish Stout, and its most famous variant, Guinness. Dark Munich malt is the most central ingredient for the local Dunkel dark lager. Vienna malt is unsurprisingly the most central ingredient in Vienna Lager.

Part of what makes these styles so easily identifiable is that their central ingredients are either not used or are of much lower importance in other styles of beer. It is helpful to visualize the distribution of eigenvector centrality for a given key ingredient node across style networks. Figure 7 shows the probability distribution of eigenvector centrality for the two most central nodes in the American IPA and Kölsch: Citra hops and Pilsner malt, respectively. While these distributions are different, they both reveal a bi-modal pattern indicating that while an ingredient may be used in many recipes, it is highly relevant in just a few. Indeed, we find that Citra hops are central components of most American ales but are missing from many European lagers. By contrast, Pilsner is the preferred base malt for many continental lagers from the Bavaria and Bohemia regions but is not as common in English and American ales, which tend to use pale ale malts such as Maris Otter Pale or Pale 2-Row as their base malt.

Figure 7: Eigenvector Centrality Distribution across Styles

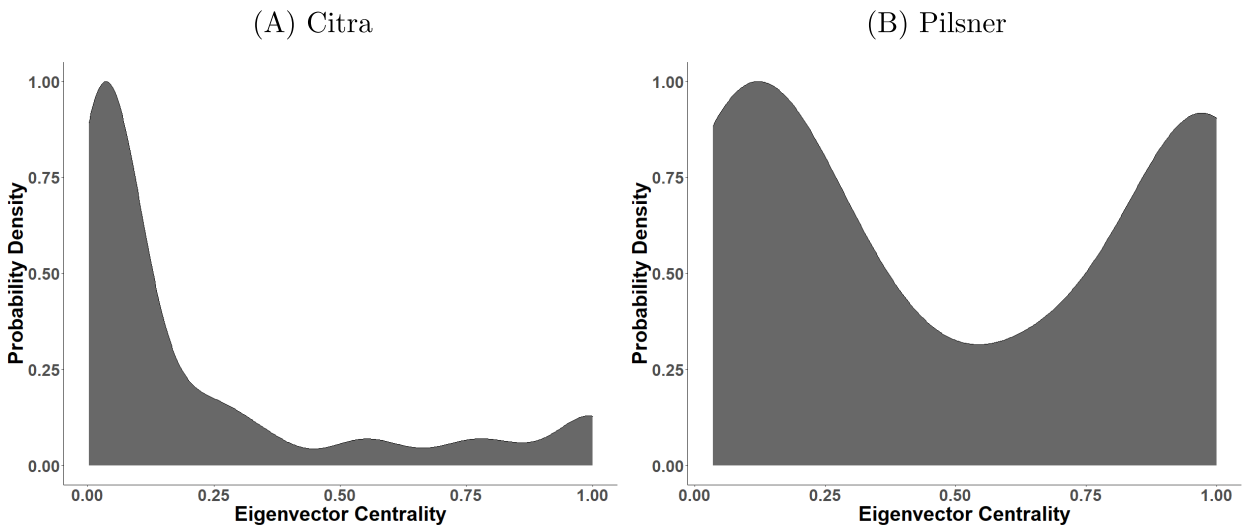

Note: These plots show the probability density function of the eigenvector centrality for two ingredients prevalent in many beer styles: Citra hops and Pilsner malt. The x-axis is the eigenvector centrality of the nodes computed for every style in our sample. The y-axis is the probability density of the centrality score. Both axes are scaled between one and zero such that when an ingredient is the most influential in a network, it will have a centrality score of one. A probability density of one means this is the most frequent centrality score of the ingredient among the beer styles. We measure eigenvector centrality according to Bonacich (1972).

To further understand how geography shapes differences in ingredient centrality it is helpful to think about two examples. Figure 8A shows the eigenvector scores of the top ingredients in two members of the same style family, American and English IPA. Considering just the hops shown in Panel 8A, it is clear that English IPA makes heavy use of American hops. Despite this colonial influence, English IPA also relies heavily on two distinctively English hops, East Kent Golding and Fuggles. These hops are conspicuously absent from American IPA, and their inclusion contributes to English IPA's unique characteristics and flavour. The American hops are bittering hops with high levels of alpha acids and citric flavor, while their English counterparts are mixed purpose hops with fewer alpha acids and are known for their earthy tones (Healey, 2016; BarthHaas, 2018). English IPA also uses Maris Otter malt, a classic malt produced in England, much more than the American IPA. Looking at the centrality scores of ingredients across the two IPAs, it is easy to see that English IPA has more herbal tones, which captures why the English version "has less hop intensity and more pronounced malt flavours than typical American versions" (BJCP, 2015).

Having considered regional differences in similar styles above, we turn to style differences within the same region. Figure 8B shows the same relationship for two German beers: Munich Helles, a light lager, and Kölsch, a pale ale. Despite belonging to two distinct style families, there is significant overlap in the centrality of their ingredients. Kölsch is the only pale ale brewed in Germany, which makes it distinct from all other beers in the country and unique to Köln. Yet, compared to ales from other nations, Kölsch uses significantly more of the base malts usually found in German pilsners and lagers. Further, Kölsch favors using the German and Czech hops abundant in lagers and known for their aroma, low bitterness, and lightly flowery and spice taste (Healey, 2016; BarthHaas, 2018). These central German nodes contribute to the uniqueness of Kölsch, a pale ale with pronounced lager traits, which could easily lead the "untrained taster to mistake it for a somewhat subtle Pils" (BJCP, 2015).

Another way to consider how beer styles differ with geography is to compare similar style networks. Along these lines, we measure the product-moment correlation coefficients between every style adjacency matrix. The correlation coefficient captures how similar the weighted edges between ingredients are across any two styles. Correlation gives us the overlap between style networks where ingredients appear frequently together and combine in similar ways. From our example in Figure 8, we ought to expect

Figure 8: Eigenvector Centrality of Ingredients by Styles

(A) American IPA vs. English IPA

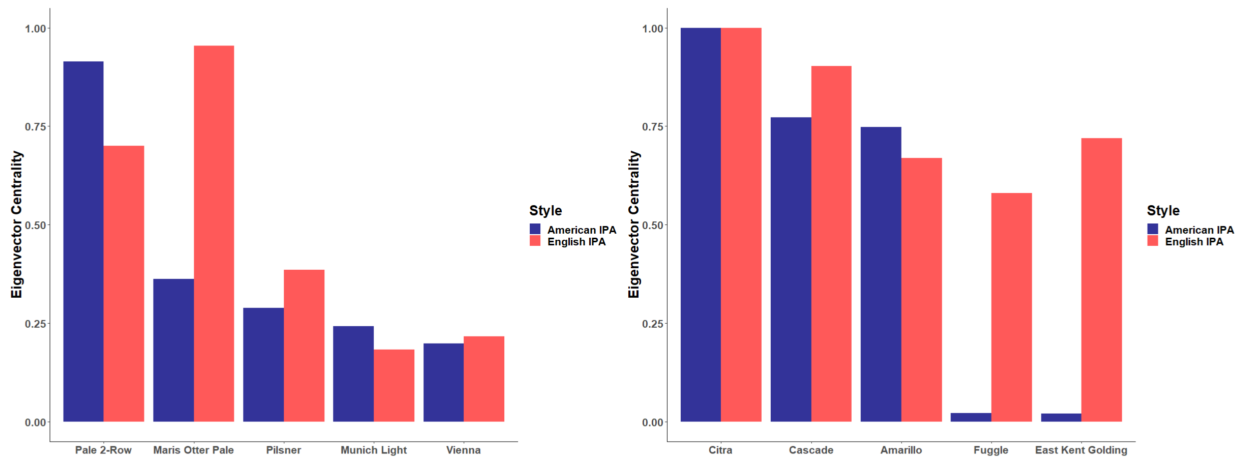

(B) Kölsch vs. Munich Helles

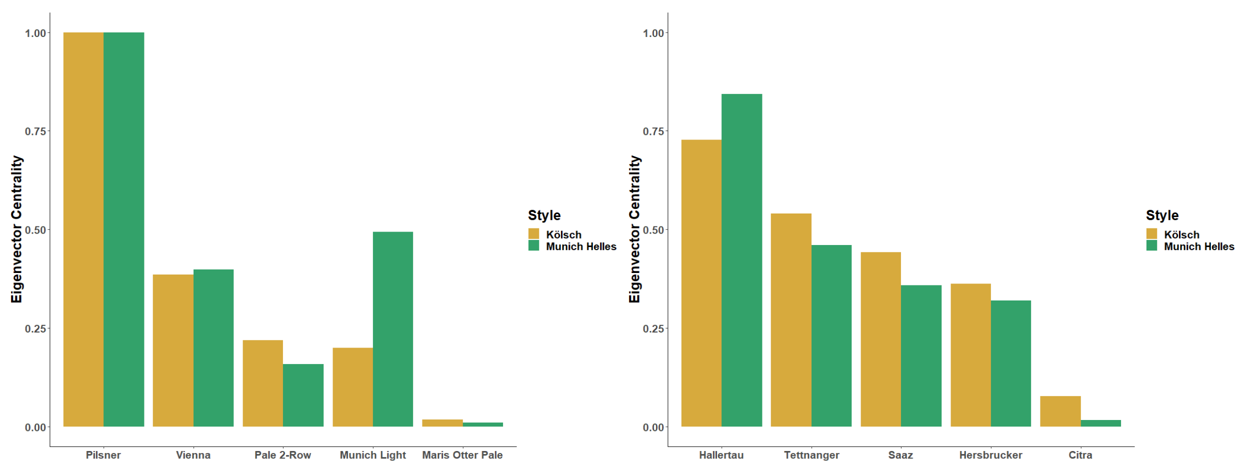

Note: These figures plot the eigenvector centrality of the five most central malts and hops in four style networks: American IPA, English IPA, Kölsch, and Munich Helles. Eigenvector centrality measures the importance of each ingredient to a style, which we compute according to Bonacich (1972). Panel 8A shows the comparison between two styles of the same family (IPA) across different countries: the United States and England. Panel 8B compares the centrality scores for two styles of different families, pale ale and pale lager, within the same country of origin: Germany. We arrange the ingredients in Panel 8A according to their centrality scores for American IPA. In Panel 8B, we arrange the ingredients according to their centrality score for Kölsch.

a higher correlation coefficient between the two German styles than their American counterpart.

Mathematically, we can express the styles correlation coefficient as:

$$cor(S, S') = \frac{cov(S, S')}{\sqrt{cov(S, S)cov(S', S')}} \tag{7}$$

where S and S' are two example adjacency matrices and their covariance is given by:

$$cov(S, S') = \frac{1}{|V_2|} \sum_{i,j} (S_{ij} - \mu_S)(S'_{ij} - \mu'_S) \tag{8}$$

where S_{ij} and S'_{ij} are the elements within each adjacency matrix, or the weighted edges between the ingredients i and j in both matrices, μ_S and μ'_S are the average degree, and $|V_2|$ is the variance. If two adjacency matrices have comparable weighted edges between their ingredients then those styles are similar and will have larger correlation coefficients.

Figure 9 lists the top ten correlations for two networks: American IPA and Kölsch. Perhaps unsurprisingly, we find American IPA to be very similar to other American beers, including the American Light Lager, particularly due to the pronounced use of American hops. Kölsch, on the other hand, is most similar to German and Bohemian lagers, and to a lesser degree to other pale ales from Europe, especially those in Belgium.

Figure 9 highlights that beer recipes and styles are clustered in space. Beer styles are more similar to other styles from the same region, even if those styles belong to very different families. This is true for our American IPA and Kölsch networks, and for other styles in different regions. For example, Bohemian Pilsner is more closely related to its regional neighbor, Czech Pale Lager (correlation coefficient of 0.95), than it is to a beer of its same style, German Pils (correlation coefficient of 0.76). Likewise, Saison, a beer style from French-speaking Wallonia in Belgium, is more similar to other Belgian ales like the Belgian Golden Ale (0.78). Golden Ale in turn is more similar to other styles from Dutch-Speaking Flanders like Belgian Golden Strong Ale (0.91) and Belgian Trippel (0.91). Figure E (Appendix) plots the correlation coefficients for all styles.

Therefore, regional ingredients are not only critical to the uniqueness and resilience

Figure 9: Top Ten Similar Styles

(A) American IPA

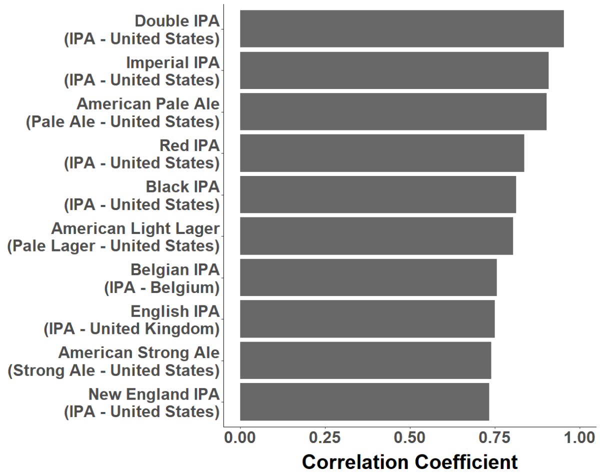

(B) Kölsch

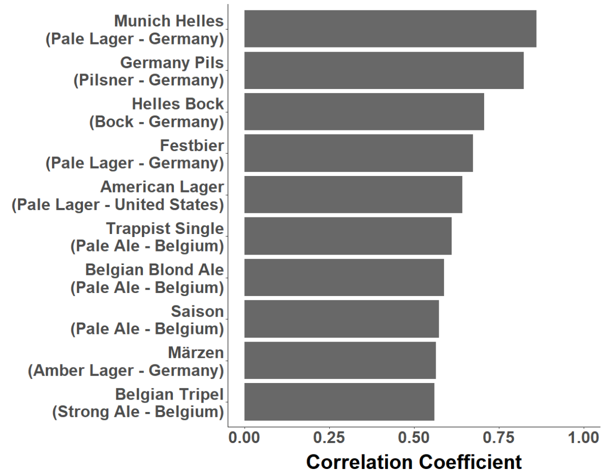

Note: These plots show the ten styles most similar to American IPA and Kölsch, ranked by correlation coefficient. The x-axis is the correlation coefficient. The y-axis displays the names of the most similar styles, with style family and country of origin in parenthesis. We calculate the correlation coefficient as product-moment correlations between any two styles' adjacency matrices.

of a style, they also transcend style boundaries and link geographically proximate beers together. This makes good sense, as the original brewers primarily had access to local ingredients and made the most with what was available. This lack of variety, be it natural or imposed, as under the German Reinheitsgebot, informed the development of these Classical styles. Even in an era of globalization, these differences persist. New World styles like American IPA benefit from the abundance of ingredients available to them. This results in a large number of ingredients (average relatedness) with a substantial number of ready substitutes (related variety). These factors give New World styles incredible resilience to losing keystone ingredients, as well as the flexibility to adapt and embrace new ones. This adaptability explains the extreme popularity of these styles and why so many brewers are drawn to them.

8 Conclusion

We bring new data and methods to the discussion on beer and place. We find that only a few key ingredients differentiate beer styles, and that geography and the diversity of ingredients matter to the resilience of a style. We are the first to collect and disambiguate a comprehensive set of beer recipe data, which we hope others will build on. Not only can this beer data answer other longstanding questions in the geography of beer literature, but the highly detailed ingredient information can also be seen as data on intermediate goods used to produce a final product. Because of this, we are able to bring an existing methodology to a new area of inquiry. We hope our use of techniques from seemingly unrelated fields inspires others to do the same. We quantify the benefits of styles having an abundance of ingredients and substitutes in their regions. This conclusion is a sensible one, and is by no means specific to beer alone. Especially in today's ever more connected world, embracing the abundance and diversity that globalization offers is useful for everyone, brewers included. We invite you to pour yourself a cold one and enjoy a sip of that diversity with us. Cheers.

Appendices

Table A: Glossary

Term	Definition	Formula
Node	Connection point in a network graph	–
Edge	Link between two nodes in a network graph	–
Edge List	Dataframe containing the starting point and end of every edge, as well as its weight	–
Adjacency Matrix	A square matrix were each element designate the edges between a pair of nodes	$S = [s_{ij}]$
Network	Graphical representation of nodes and edges	–
Degree (Centrality)	Sum of all the edges incident to a node	$\sum_{i \neq j}^{n} s_{ij}$
Eigenvector Centrality	A measure of a node's influence. We calculate it as the weighted sum of the centrality of all adjacent nodes	$c(v_i) = \sum_{j=1}^{n} s_{ij} c(v_j)$
Density	The share of existing edges out of all possible links in the network	$\frac{m}{n(n-1)/2}$
Diameter	The largest distance between any two pair of nodes	–
Clustering Coefficient	The proportion of exiting edges among each node's neighbors	–
Resiliency	Percentage of nodes one can delete before the network becomes fragmented into many unconnected components	–
Molloy-Reed Criterion	Threshold at which a complex network will lose its large connected component	$\Omega_s = \frac{\sum_{i=1}^{N} k_{is}^2}{\sum_{i=1}^{N} k_{is}} < 2$
Unrelated Variety	A measure of diversity among components of a recipe-ingredient incidence matrix	$\sum_{i=1}^{I} P_{is} \log_2 \left(\frac{1}{P_{is}} \right)$
Relatedness	A measure of similarity between the networks nodes in the global network	$\frac{s_{ij}}{\sqrt{N_i * N_j}}$
Average Relatedness	The average relatedness across all nodes in a style's network	$\frac{\sum_i \sum_j R_{ij}(N_i N_j) + \sum_i 2N_i}{P_{is}(P_{is}-1)}$

Note: The table shows definitions and formulas for all the network related terms used throughout the paper. For further information on these, we refer to Wasserman et al. (1994); Barabási et al. (2016); Frenken et al. (2007); Kogler et al. (2013).

Table B: Style List

Style	Family	Country	Region	No. Recipes
American IPA	IPA	United States	West Coast	14694
American Pale Ale	Pale Ale	United States	West Coast	12363
American Light Lager	Pale Lager	United States	Midwest	4638
Saison	Pale Ale	Belgium	Wallonia	4163
Blonde Ale	Pale Ale	United States	West Coast	3605
New England IPA	IPA	United States	New England	3020
American Amber Ale	Amber Ale	United States	West Coast	2924
Irish Red Ale	Amber Ale	Ireland	Ireland	2133
American Stout	Stout	United States	West Coast	1996
Weissbier	Wheat Beer	Germany	Bavaria	1818
Witbier	Wheat Beer	Belgium	Flemish Brabant	1704
Strong Bitter	Amber Ale	United Kingdom	England	1582
Sweet Stout	Stout	United Kingdom	England	1571
American Porter	Porter	United States	East Coast	1560
English IPA	IPA	United Kingdom	England	1527
Oatmeal Stout	Stout	United Kingdom	England	1513
Imperial IPA	IPA	United States	West Coast	1498
American Brown Ale	Brown Ale	United States	West Coast	1452
Double IPA	IPA	United States	West Coast	1366
Russian Imperial Stout	Stout	Russia	Baltic	1341
Black IPA	IPA	United States	West Coast	1136
Best Bitter	Amber Ale	United Kingdom	England	1128
Ordinary Bitter	Amber Ale	United Kingdom	England	1111
British Brown Ale	Brown Ale	United Kingdom	England	1069
California Common	Amber Lager	United States	San Francisco	1015
Belgian Pale Ale	Pale Ale	Belgium	Flemish Brabant	1008
American Wheat Beer	Wheat Beer	United States	Pacific Northwest	969
Kölsch	Pale Ale	Germany	Cologne	946
Belgian Blond Ale	Pale Ale	Belgium	Flemish Brabant	918
Märzen	Amber Lager	Germany	Bavaria	910
Red IPA	IPA	United States	West Coast	893
Cream Ale	United States	United States	Midwest	857
Berliner Weisse	Wheat Beer	Germany	Berlin	824
Belgian Dubbel	Amber Ale	Belgium	Flemish Brabant, Antwerp	812
Robust Porter	Porter	American	East Coast	720
Belgian Tripel	Strong Ale	Belgium	Antwerp	712
British Golden Ale	Pale Ale	United Kingdom	England	705
American Lager	Pale Lager	United States	Midwest	682
Brown Porter	Porter	United Kingdom	England	674
Dunkles Weissbier	Wheat Beer	Germany	Bavaria	612
Irish Stout	Stout	Ireland	Ireland	583
Rye IPA	IPA	United States	West Coast	573
Belgian Golden Strong Ale	Strong Ale	Belgium	Flemish Brabant, Antwerp	565

Belgian Dark Strong Ale	Strong Ale	Belgium	Flemish Brabant, Antwerp	563
German Pils	Pilsner	Germany	Bavaria	558
White IPA	IPA	United States	West Coast	553
Imperial Stout	Stout	United Kingdom	England	546
Vienna Lager	Amber Lager	Austria	Vienna	503
Dark Mild	Brown Ale	United Kingdom	England	501
Foreign Extra Stout	Stout	United Kingdom	England	475
English Porter	Porter	United Kingdom	England	465
Bohemian Pilsener	Pilsner	Czech Republic	Plzeň	451
American Barleywine	Strong Ale	United States	West Coast	441
Scottish Export	Amber Ale	United Kingdom	Scotland	439
English Barleywine	Strong Ale	United Kingdom	England	426
Old Ale	Strong Ale	United Kingdom	England	400
Belgian IPA	IPA	Belgium	Flanders	386
British Strong Ale	Strong Ale	United Kingdom	England	384
Belgian Specialty Ale	Strong Ale	Belgium	Flemish Brabant, Antwerp	376
Dry Stout	Stout	United Kingdom	England	375
American Strong Ale	Strong Ale	United States	West Coast	359
Weizenbock	Bock	Germany	Bavaria	346
Doppelbock	Bock	Germany	Bavaria	332
Baltic Porter	Porter	Scandinavia	Baltic	331
Helles Bock	Bock	Germany	Bavaria	319
Bière de Garde	Pale Ale	France	Northern France	311
Altbier	Amber Ale	Germany	Düsseldorf	304
Mild	Amber Ale	United Kingdom	England	299
Munich Dunkel	Dark Lager	Germany	Bavaria	290
Scottish Heavy	Amber Ale	United Kingdom	Scotland	290
Trappist Single	Pale Ale	Belgium	Flemish Brabant, Antwerp	255
Strong Scotch Ale	Strong Ale	United Kingdom	Scotland	248
Flanders Red Ale	Sour Ale	Belgium	West Flanders	243
Schwarzbier	Dark Lager	Germany	Thuringia, Saxony, & Franconia	241
Munich Helles	Pale Lager	Germany	Munich	238
Wee Heavy	Strong Ale	United Kingdom	Scotland	222
Czech Premium Pale Lager	Pale Lager	Czech Republic	Plzeň	213
Scottish Light	Pale Ale	United Kingdom	Scotland	183
Lambic	Wheat Beer	Belgium	Brussels	172
Australian Sparkling Ale	Pale Ale	Australia	Australia	154
Brown IPA	IPA	United States	Westcoast	153
Traditional Bock	Bock	Germany	Einbeck	151
Roggenbier	Wheat Beer	Germany	Regensburg	150
Festbier	Pale Lager	Germany	Munich	147
Czech Pale Lager	Pale Lager	Czech Republic	Plzeň	141
Fruit Lambic	Wheat Beer	Belgium	Brussels	138

Irish Extra Stout	Stout	Ireland	Ireland	131
Gose	Wheat Beer	Germany	Leipzig	119
Gueuze	Wheat Beer	Belgium	Brussels	116
Oud Bruin	Sour Ale	Belgium	East Flanders	110

Table C: Malt List

Malt Name	Recipes	EBC	Styles	Countries
Pale 2-Row	37806	2.18	American Lager, Cream Ale, etc.	AR, AU, BE, CA, DE, DK, FI, FR, UK, US
Pilsner	25483	1.70	Altbier, Belgian Tripel, Berliner Weisse, etc.	AR, AU, BE, BR, CA, DE, FI, FR, NL, NZ, UK, US
Maris Otter Pale	19820	3.66	Best Bitter, Mild, Ordinary Bitter, etc	BE, CA, UK, US
Pale	14814	2.95	Belgian IPA	AR, AU, BE, BR, CA, CL, DE, FI, IE, NL, NZ, SE, UK, US, ZA
Wheat	13826	2.09	Dunkles Weissbier, Fruit Lambic, Weizenbock	AR, AU, BE, CA, DE, FI, IE, NL, NZ, UK, US
Chocolate	13654	371.69	American Porter, American Stout, etc	AR, AU, BE, CA, CL, DE, FI, IE, NL, NZ, UK, US
Munich Light	11486	8.07	Märzen	AU, BE, CA, DE, FI, NL, UK, US
Vienna	11322	3.95	Festbier, Vienna Lager	AR, AU, BE, BR, CA, CL, DE, FI, FR, IE, NL, NZ, UK, US
Caramel/Crystal60L	10970	59.95	American Brown Ale	AR, BE, CA, DE, NL, UK, US
Carapils Dextrine	9424	1.80	Strong Scotch Ale	DE, FI, US
Roasted Barley	9092	411.03	Irish Extra Stout	AR, AU, BE, CA, DE, FI, NL, NZ, UK, US
Caramel/Crystal40L	7790	40.01	Robust Porter	CA, DE, UK, US
CaraMunich	7682	47.03	Flanders Red Ale	BE, DE, NL, UK, US
Munich	7213	14.42	Traditional Bock	AR, AU, BE, BR, CA, CL, DE, FI, IE, NL, NZ, UK, US
CaraPils	6659	2.39	Munich Helles	AR, BE, DE, IE, UK, US
White Wheat	6654	2.75	American Wheat Beer	BE, CA, DE, US
Acidulated	5918	3.38	Gose	BE, DE
Caramel/Crystal120L	5000	120.00	Imperial Stout	BE, CA, DE, UK, US
Caramel/Crystal20L	4692	20.20	American IPA	BE, CA, DE, FI, US
SpecialB	4312	116.08	Belgian Dubbel	BE, UK
Biscuit	4072	23.21	Belgian Dubbel	AR, AU, BE, CA, DE, FI, NL, NZ, UK, US
Dark Munich	4042	15.71	Munich Dunkel	AU, CA, DE, FI, NL, UK, US
Rye	3992	3.56	Roggenbier, Rye IPA	AU, BE, CA, DE, FI, NZ, UK, US
Honey	3828	24.60	Scottish Light	CA, FI, UK, US
Victory	3784	27.98	American Brown Ale	UK, US

Golden Promise	3750	2.96	Wee Heavy	FI, IE, UK
Melanoidin	3530	24.87	Traditional Bock	AR, BE, DE, US
Aromatic	3332	29.24	Flanders Red Ale	BE, DE, FR, NL, UK, US
Caramel/Crystal10L	3107	10.02	Blonde Ale	BE, DE, FI, NL, US
Caramel/Crystal80L	3027	80.01	American Brown Ale	CA, UK, US
Black	2690	506.76	Robust Porter	BE, DE, FI, NL, UK, US
Pale Chocolate	2517	209.35	English Porter	FI, NZ, UK, US
CaraRed	2311	20.00	Red IPA	DE, UK, US
Black Patent	2101	523.74	Robust Porter	UK, US
CarafaIII	2057	533.70	Black IPA	DE, UK
CaraHell	2036	10.91	Helles Bock	DE, US
Caramel/Crystal30L	2029	29.84	Scottish Export	AR, FI, UK, US
Brown	2017	69.07	English Porter	NL, NZ, UK, US
Pale Wheat	1965	1.71	Weizenbock	BE, CA, DE, UK, US
Caramalt	1930	18.11	Irish Extra Stout	AU, BE, CL, DE, UK, US
Carafoam	1787	1.89	American Light Lager	DE, US
Caramel/Crystal15L	1747	15.00	American Strong Ale	AR, UK, US
CarafaII	1735	426.81	Schwarzbier	DE
Caramel/Crystal45L	1621	44.15	Irish Extra Stout	BE, NL, UK, US
CaraAroma	1609	126.07	Foreign Extra Stout	DE, UK, US
Bohemian Pilsner	1544	1.92	Bohemian Pilsner	DE
Dark Crystal	1499	88.99	English Porter	AU, NZ, UK, US
Amber	1494	27.19	English Porter	AU, BE, CA, DE, IE, NL, UK, US
Caramel/Crystal90l	1395	90.01	Foreign Extra Stout	CL, UK, US
Midnight Wheat	1347	549.76	Black IPA	CA, US
Red Wheat	1281	2.48	Weizenbock	CA, US
Ale	1258	2.96	British Golden Ale	AU, BE, DE, FI, IE, NL, NZ, UK, US
Oats	1257	2.25	New England IPA	AU, BE, CA, FI, IE, NZ, UK, US
Smoked	1207	4.11	Strong Scotch Ale	BE, DE, FI, NL, NZ, UK, US
Caramel/Crystal150L	1150	125.26	Oatmeal Stout	BE, DE, FI, UK, US
CarafaI	1126	340.71	Irish Extra Stout	DE
Extra Dark Crystal	1096	133.84	Scottish Light	AU, UK
CaraVienne	1063	20.22	Oud Bruin	BE, US
CaraAmber	1025	23.55	American Amber Ale	DE
Pale 6-Row	1020	1.82	Cream Ale	CA, US
Golden Naked Oats	987	9.79	American Brown Ale	UK
Special Roast	963	50.14	American Brown Ale	BE, US
Chocolate Dark	935	459.39	Irish Extra Stout	FI, NZ, US
Caramel/Crystal50L	857	43.72	Foreign Extra Stout	DE, FI, UK
Caramel/Crystal75L	779	75.00	Brown Porter	US
RedX	721	11.98	Red IPA	DE, US

Chocolate Wheat	647	411.66	Irish Extra Stout	DE, NL, US
Caramel/Crystal70l	625	69.22	English Porter	FI, UK
Chocolate Rye	606	243.60	Irish Extra Stout	DE, NZ, US
Caramel/Crystal140L	585	139.75	English Porter	BE, CL, DE, UK, US
Blackprinz	584	500.00	Schwarzbier	US
Crystal Light	566	36.06	Brown Porter	AU, NZ, UK
Floor Malted Bohemian Pilsner	545	2.15	Czech Premium Pale Lager	DE
Coffee	531	158.31	English Porter	BE, CA, CL, DE, FR, NL, NZ, UK, US
Dextrine	524	1.75	Brown Porter	FI, UK, US
Flaked Oats	478	2.07	Oatmeal Stout	BE, CA, DE, NL, UK, US
Abbey	475	17.18	Belgian Dubbel	BE, DE, US
Roasted Wheat	475	49.37	Fruit Lambic	AU, BE, FI, NZ, UK, US
Lager	474	1.53	British Golden Ale	IE, NL, UK
Crystal Medium	472	57.59	English Porter	AU, NZ, UK
Black Barley	469	525.93	Oatmeal Stout	UK, US
Carastan	436	33.80	Russian Imperial Stout	UK
Extra Pale	429	7.76	Belgian Specialty Ale	AU, BE, CL, DE, FI, UK, US
Darkwheat	420	6.55	Weizenbock	DE
CaraCrystal Wheat	404	54.67	Brown IPA	BE, US
Debittered Black	393	566.10	Imperial Stout	BE
Crystal Rye	380	88.59	Black IPA	NZ, UK
Dehusked CarafIII	336	470.22	Black IPA	DE
Carabrown	333	55.00	American Brown Ale	US
Melanoid	307	27.54	Traditional Bock	BE, DE, FI, NL
Pearl	299	3.28	Strong Scotch Ale	UK, US
Spelt	297	2.04	Lambic	DE, NL
CaraPale	290	3.89	Irish Extra Stout	BE, FI, UK, US
Caramel/Crystal100L	288	40.22	Irish Extra Stout	FI, UK, US
Peated	287	2.51	Strong Scotch Ale	BE, DE, UK
CaraWheat	283	46.05	Dunkles Weissbier	DE, FI
Flaked Barley	283	1.62	Irish Stout	DE, US
Carabelge	257	13.51	Flanders Red Ale	BE, DE
Caramel/Crystal55L	241	55.09	Oud Bruin	CL, UK
CaraBohemian	238	75.00	Robust Porter	DE
Crystal	232	61.44	Brown Porter	AU, BE, IE, NZ, UK, US
Dehusked CarafII	229	418.03	Schwarzbier	DE
CaraRye	228	66.19	Black IPA	DE, NZ, US
ESB	209	3.47	Mild	CA, UK
Toffee	201	5.65	Irish Extra Stout	NZ
CaraGold	180	33.47	Festbier	BE, IE, UK, US
Gladiator	176	5.14	Fruit Lambic	NZ
Chocolate Light	169	419.60	Irish Extra Stout	AU, FI, NZ, UK, US
Flaked Wheat	164	2.03	Witbier	BE, CA, DE, NZ, US

Mild	136	3.11	Mild	UK
Ashburne Mild	131	5.30	Dark Mild	US
Aurora	125	28.43	Belgian Pale Ale, Brown IPA	NZ
Cookie	115	19.26	British Golden Ale	FI
Dehusked CarafI	115	337.29	Schwarzbier	DE
Cararuby	109	21.12	Red IPA	BE, UK
Caramel/Crystal300L	108	114.82	Traditional Bock	FI
Redback	106	33.86	Red IPA	NL, NZ
Shepherds Delight	103	144.93	Red IPA	NZ
Floor Malted Bohemian Wheat	101	1.98	Gueuze	DE
Sourgrapes	101	2.18	Gose	NZ
Crystalmillet	96	17.67	Helles Bock	US
Kölsch	92	3.29	Kölsch	DE
Buckwheat	91	2.30	Irish Stout	US
Carastan Light	91	15.00	Scottish Export	UK
Pale Millet	89	1.65	Old Ale	US
Stout	89	18.78	Irish Stout	IE
Chit	85	1.52	Scottish Light	DE, NZ, UK, US
Double Roast Crystal	79	110.04	Scottish Heavy	UK, US
SpecialW	79	114.72	English Porter	BE, DE
Caraplus150	78	59.20	Foreign Extra Stout	FI
Rice	72	58.19	American Wheat Beer	US
Optic	70	2.15	Strong Scotch Ale	UK
Caraplus100	66	38.00	Black IPA	FI
CaraBlond	63	7.99	White IPA	BE
Caramel/Crystal250L	63	94.00	Roggenbier	FI
Dark	58	71.12	Brown IPA	CA, DE, FI, NZ, US
Supernova	57	57.62	Brown IPA	NZ
Rolled Oats	56	1.73	Scottish Heavy	NZ
Flakedrice	54	0.93	American Light Lager	US
Pale Crystal	52	28.92	British Strong Ale	UK
Pale Rye	50	3.29	Roggenbier	DE, UK, US
Red	50	22.05	Red IPA	DE, FI, NL, US
Flaked Rye	49	2.29	Dry Stout	US
Munich Millet	49	6.61	Scottish Light	US
Black Pearl	46	350.43	Oatmeal Stout	CL, FI, NZ, US
Heritage	45	6.35	Old Ale	CA, UK
Flaked Corn	42	0.71	Cream Ale	US
Roasted Oats	40	6.22	Belgian Dubbel	CA, UK, US
Dark Roasted Millet	39	298.21	English Porter	US
Caraplus200	38	76.67	Russian Imperial Stout	FI
Specialx	38	128.09	Brown Porter	DE
Distillers	37	2.37	Gueuze	IE, NZ, US
Smoked Wheat	35	2.31	Roggenbier	DE, US
Caramel/Crystal200l	32	83.38	British Strong Ale	BE, FI, UK, US

Lager Light	32	1.43	Fruit Lambic	NZ
Fullpint	30	3.45	Strong Bitter	US
Heidelberg	26	1.34	Munich Helles	DE
Opal 44	24	35.33	American Amber Ale	US
CaraMillet	22	32.32	Irish Red Ale	US
Enzyme	21	0.40	Red IPA	FI
Pale Compass	21	2.69	Old Ale	AU
Red Crystal	21	98.59	Foreign Extra Stout	UK
Light Roasted Millet	20	6.70	Irish Red Ale	US
Base	19	1.97	Rye IPA	NZ, US
Caramel/Crystal400L	19	198.24	British Strong Ale	CA, DE, FI, UK
Caramel/Crystal65L	19	65.00	Scottish Light	UK, US
Caraplus250	19	97.26	American Light Lager	FI
Lamonta	19	2.89	Belgian IPA	US
Barley	18	195.82	Imperial Stout	BR, DE, NL, NZ
Golden	18	4.07	Gose	FI, US

Table D: Hop List

Hop Name	Recipes	Alpha	Styles	Countries
Cascade	18582	6.91	American Barleywine, American Pale Ale, etc.	United States
Citra	15950	11.88	American IPA, Double IPA, Black IPA, etc.	United States
Centennial	11046	9.79	American Barleywine	United States
Amarillo	10867	8.60	Belgian IPA	United States
East Kent Golding	9833	5.15	British Strong Ale, Scottish Export, Strong Bitter, etc.	United Kingdom
Magnum	9696	13.90	Altbier	Germany
Simcoe	9223	12.96	Imperial IPA	United States
Mosaic	9151	12.58	New England IPA	United States
Columbus	8245	15.01	Imperial IPA	United States
Chinook	7568	12.68	American Barleywine	United States
Saaz	6772	3.48	Bohemian Pilsner, Czech Premium Lager, etc.	Czech Republic
Fuggle	6538	4.57	Mild	United Kingdom
Hallertau	5328	3.88	Munich Helles	Germany
Willamette	4786	4.69	American Brown Ale	United States
Galaxy	3995	14.44	New England IPA	Australia
Northernbrewer	3926	7.95	California Common	United States
Warrior	3127	15.90	Double IPA	United States
Nugget	3052	13.66	Old Ale	United States
Tettnanger	2986	4.18	Altbier	Germany
Styrian Gold	2816	4.87	Belgian Golden Strong Ale	Slovenia
Perle	2721	7.67	Doppelbock	Germany
Hersbrucker	2572	3.68	Munich Dunkel	Germany
El Dorado	2240	15.06	New England IPA	United States
Challenger	2221	7.91	Strong Bitter	United Kingdom
Nelson Sauvin	2159	12.17	Brown IPA	New Zealand
Golding	1671	4.83	British Strong Ale	United States
Azacca	1549	13.75	New England IPA	United States
Mandarina Bavaria	1330	8.33	White IPA	Germany
Summit	1225	17.28	Black IPA	United States
Target	1207	10.78	Strong Bitter	United Kingdom
Motueka	1193	6.84	Belgian IPA	New Zealand
Crystal	1158	4.18	Brown Porter, Cream Ale	United States
Mounthood	1105	5.06	Cream Ale	United States
Galena	1076	13.00	Fruit Lambic	United States
Ekuanot	949	15.20	New England IPA	United States
Hallertau Blanc	912	9.40	Gose	Germany
Cluster	910	6.87	Cream Ale	United States
Sorachi Ace	887	11.46	Saison	Japan
Apollo	814	18.98	Imperial IPA	United States
Ahtanum	788	5.50	Brown IPA	United States

Bravo	760	15.19	White IPA	United States
Falconer's Flight	738	10.16	Black IPA	United States
Liberty	725	4.19	Cream Ale	United States
Lemondrop	712	5.80	Berliner Weisse	United States
Hullmelon	693	6.75	Saison	Germany
Vic Secret	624	16.34	New England IPA	Australia
Sterling	571	8.42	Doppelbock	United States
Bramling Cross	545	6.40	Old Ale	United Kingdom
Brewer's Gold	524	8.28	Belgian Dark Strong Ale	United Kingdom
Horizon	518	12.19	American Amber Ale	United States
Spalt	510	4.17	Altbier	Germany
Idaho 7	503	13.08	New England IPA	United States
Hallertauer Tradition	481	5.43	Festbier	Germany
Calypso	466	13.08	Fruit Lambic	United States
Glacier	463	5.42	Irish Stout	United States
First Gold	444	7.84	British Golden Ale	United Kingdom
Comet	397	10.48	Brown IPA	United States
Pacific Gem	385	14.90	White IPA	New Zealand
Zeus	377	16.13	Belgian IPA	United States
Northdown	372	8.20	Best Bitter	United Kingdom
Admiral	346	14.22	British Golden Ale	United Kingdom
Pacific Jade	345	13.26	Brown IPA	New Zealand
Loral	340	11.68	Fruit Lambic	United States
Rakau	325	10.44	Red IPA	New Zealand
Sabro	323	14.61	New England IPA	United States
Strisselspalt	292	3.98	Lambic	France
Zythos	286	10.72	Belgian Specialty Ale	United States
Cashmere	279	7.96	Gueuze	United States
Palisade	278	7.66	American Strong Ale	United States
Aurora	275	8.21	Belgian Pale Ale, Brown IPA	Slovenia
Equinox	265	13.83	Brown IPA	United States
Belma	262	10.26	American Barleywine	United States
Wakatu	256	7.36	Irish Extra Stout	New Zealand
Ella	247	14.72	British Golden Ale	Australia
Waimea	240	16.55	Belgian IPA	New Zealand
Denali	235	14.12	Belgian IPA	United States
Waiiti	234	2.97	Lambic	New Zealand
Herkules	229	15.33	Weizenbock	Germany
Pacifica	226	5.14	Czech Pale Lager	New Zealand
Tomahawk	223	14.79	Brown IPA	United States
Saphir	214	3.94	Festbier	Germany
Jarrylo	201	14.84	Rye IPA	United States
Topaz	199	15.76	White IPA	Australia
Southern Cross	194	12.74	Old Ale	New Zealand
Greenbullet	188	11.71	Red IPA	New Zealand
Pilgrim	184	10.26	British Golden Ale	United Kingdom
Kohatu	163	6.62	Belgian IPA	New Zealand

Progress	159	6.29	Best Bitter	United Kingdom
Southern Star	158	11.28	British Golden Ale	South Africa
Dr.Rudi	156	11.43	Berliner Weisse	New Zealand
Summer	155	6.10	Fruit Lambic	Australia
Polaris	152	19.19	Brown IPA	Germany
Pekko	144	15.09	Irish Extra Stout	United States
Enigma	143	15.77	Fruit Lambic	Australia
Eureka	139	16.82	British Golden Ale	United States
Riwaka	138	5.28	Belgian IPA	New Zealand
Phoenix	135	10.52	Irish Extra Stout	United Kingdom
Millennium	132	15.67	English Barleywine	United States
Celeia	116	3.57	Trappist Single	Slovenia
Strata	116	12.77	New England IPA	United States
Vanguard	114	5.12	Vienna Lager	United States
Spalterselect	110	4.36	Festbier	Germany
Medusa	104	3.96	New England IPA	United States
Pride of Ringwood	103	9.76	American Lager	Australia
Bullion	98	8.36	Scottish Heavy	United Kingdom
Columbia	91	10.38	Traditional Bock	United States
Moutere	85	17.72	Belgian IPA	New Zealand
African Queen	84	12.14	British Golden Ale	South Africa
Bru1	84	14.41	New England IPA	United States
Santiam	79	6.62	Gose	United States
Southern Promise	77	11.46	Brown IPA	Slovenia
Aramis	76	7.17	Saison	France
Bobek	76	4.84	Belgian Tripel	Slovenia
Opal	74	6.88	Gueuze	Germany
Taiheke	68	7.16	British Strong Ale	New Zealand
Archer	66	4.43	Old Ale	United Kingdom
Premiant	64	8.91	Czech Pale Lager	Czech Republic
Lublin	63	4.39	Roggenbier	Poland
Meridian	63	7.00	Red IPA	United States
Kazbek	62	6.00	Munich Dunkel	Czech Republic
Marynka	62	8.53	Belgian IPA	Poland
Sladek	61	6.40	Czech Premium Pale Lager	Czech Republic
Newport	60	13.44	Scottish Light	United States
Jester	59	8.57	British Golden Ale	United Kingdom
Styrian Wolf	52	12.66	English Barleywine	Slovenia
Sovereign	51	5.42	British Golden Ale	United Kingdom
Pioneer	47	8.87	British Golden Ale	United Kingdom
Callista	45	3.79	Weizenbock	Germany
Hallertauer	42	4.56	Festbier	Germany
Superalpha	42	12.75	English IPA	New Zealand
Triskel	41	4.86	Fruit Lambic	France
Ariana	38	10.60	Belgian Specialty Ale	Germany
Mount Rainier	38	6.18	Munich Dunkel	United States
HBC472	37	9.65	Cream Ale	United States

Zappa	37	6.38	New England IPA	United States
Sticklebract	31	11.76	American Light Lager	New Zealand
Superpride	31	13.80	British Golden Ale	Australia
Pahto	30	17.07	Belgian Specialty Ale	United States
Helga	29	5.93	Gose	Australia
Sonnet	29	5.22	Strong Scotch Ale	United States
Boadicea	28	7.01	British Golden Ale	United Kingdom
HBC438	27	15.66	White IPA	United States
Olicana	27	7.89	English IPA	United Kingdom
Dana	26	11.81	Old Ale	Slovenia
Endeavour	26	9.07	Mild	United Kingdom
HBC342	26	11.94	Imperial IPA	United States
Julius	26	7.36	English IPA	United States
Smaragd	25	5.67	Weizenbock	Germany
Ultra	25	5.09	American Lager	United States
Caliente	24	14.22	American Barleywine	United States
Legacy	24	7.75	British Golden Ale	United States
Minstrel	23	6.05	Weizenbock	United Kingdom
Delta	22	6.39	British Strong Ale	United States
Sybilla	20	5.67	Fruit Lambic	Poland
Merkur	19	12.85	Schwarzbier	Germany
Michigan Copper	19	9.32	Rye IPA	United States
Tahoma	19	6.69	Trappist Single	United States
Triplepearl	19	9.55	Helles Bock	United States
Hallertauer Taurus	17	15.50	Altbier	Germany
Monroe	17	3.36	American Strong Ale	Germany
Lubelska	16	4.39	Belgian Specialty Ale	Poland

Figure A: An Abridged Sample BeerXML File

```xml
<?xml version="1.0" encoding="UTF-8"?>
 <RECIPE>
  <NAME>Avg. Perfect Northeast IPA (NEIPA)</NAME>
  <VERSION>1</VERSION>
  <TYPE>All Grain</TYPE>
     ...
  <FERMENTABLES>
   <FERMENTABLE>
    <NAME>Pale 2-Row</NAME>
    <TYPE>Grain</TYPE>
    <AMOUNT>4.8761179775</AMOUNT>
    <YIELD>80.43</YIELD>
    <COLOR>1.8</COLOR>
   </FERMENTABLE>
     ...
  </FERMENTABLES>
  <HOPS>
   <HOP>
    <NAME>Citra</NAME>
    <ALPHA>12.6</ALPHA>
    <AMOUNT>0.0283495231</AMOUNT>
    <USE>Boil</USE>
    <USER_HOP_USE>Boil</USER_HOP_USE>
    <TIME>10</TIME>
    <FORM>Pellet</FORM>
   </HOP>
     ...
  </HOPS>
  <MISCS>
   <MISC>
    <NAME>Irish Moss</NAME>
    <TYPE>Fining</TYPE>
    <USE>Boil</USE>
    <TIME>15</TIME>
    <AMOUNT>0.00246446</AMOUNT>
   </MISC>
     ...
  </MISCS>
     ...
  <STYLE>
   <NAME>Specialty IPA: New England IPA</NAME>
   <CATEGORY>IPA</CATEGORY>
   <CATEGORY_NUMBER>21</CATEGORY_NUMBER>
   <STYLE_LETTER>B</STYLE_LETTER>
   <STYLE_GUIDE>BJCP</STYLE_GUIDE>
   <TYPE>Ale</TYPE>
   ...
  </STYLE>
 </RECIPE>
```

Figure B: Stress Test - Kölsch

(A) Targeted Attacks

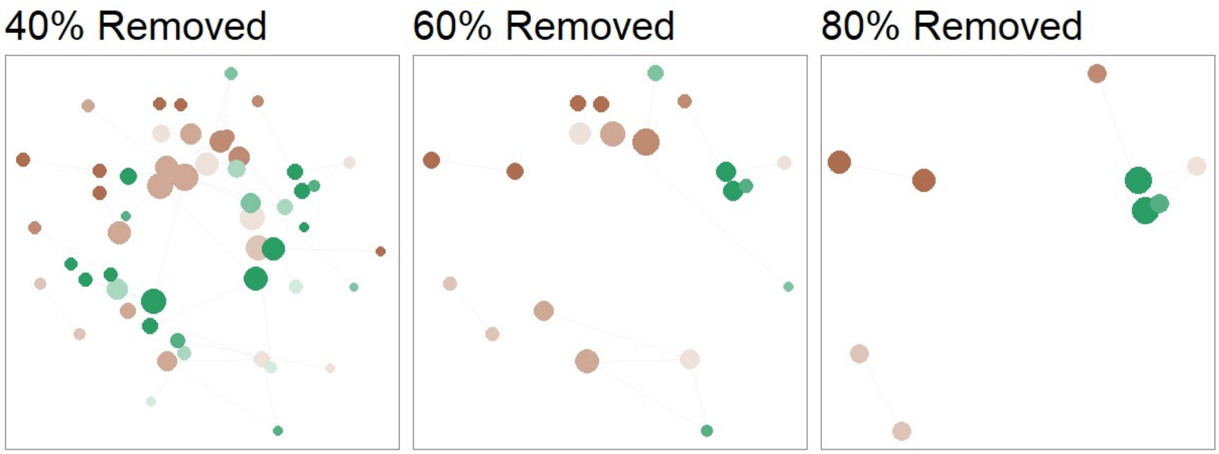

(B) Random Attacks

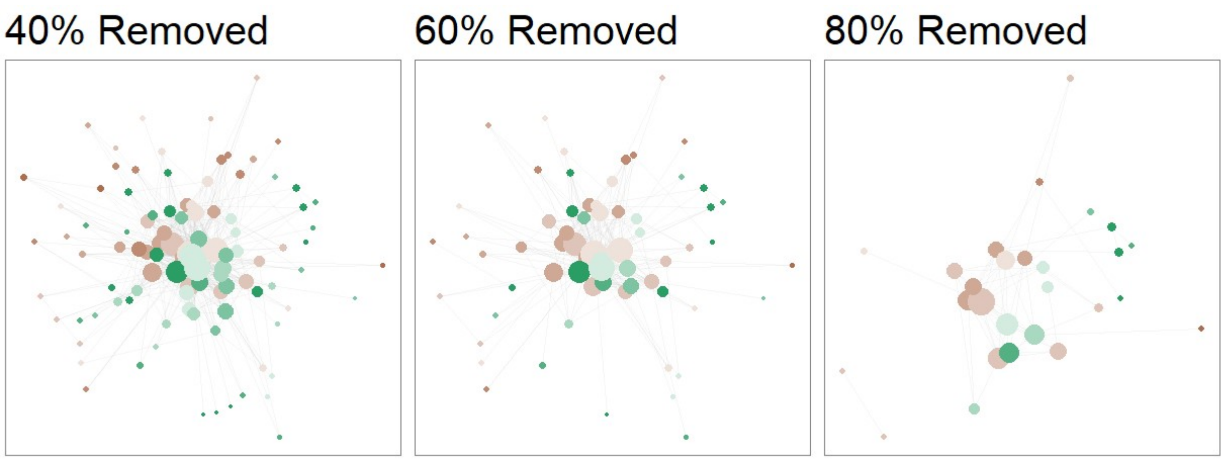

Note: Both panels depict the impact of removing 40%, 60% or 80% of the nodes from the Kölsch network. Panel BA shows the effect of targeted deletion according to eigenvector centrality. Panel BB shows a random attack where nodes are deleted in random order.

Figure C: Resiliency Against Random Removal - Kölsch

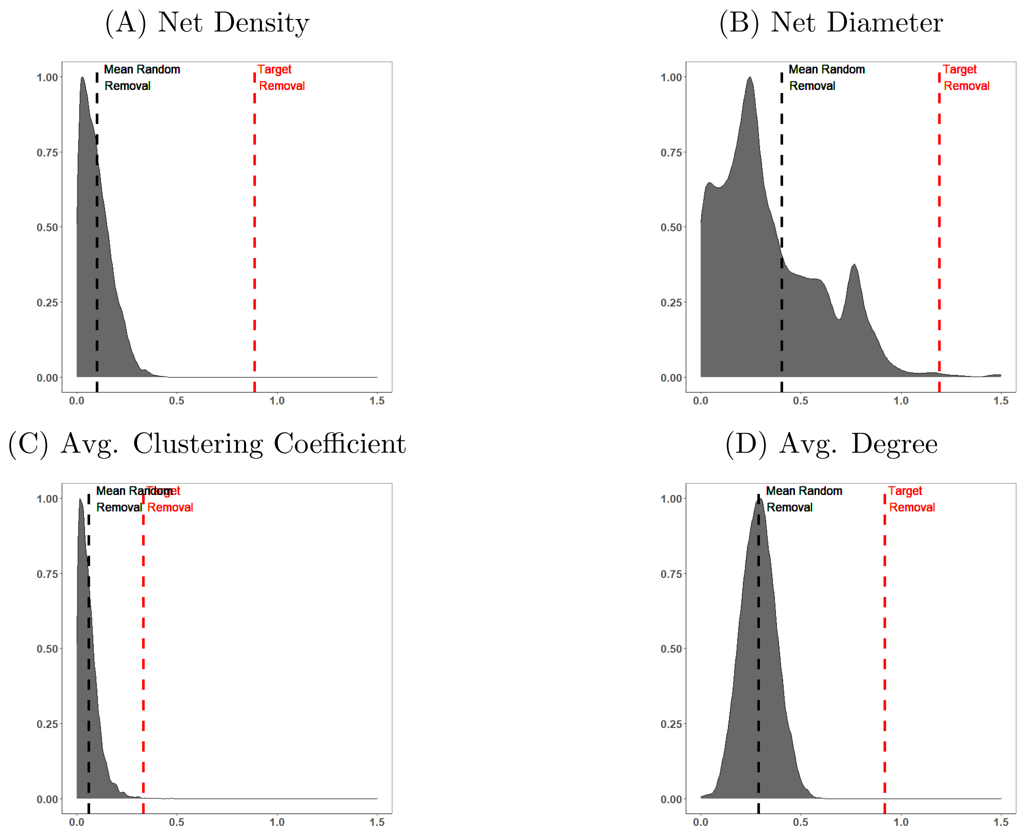

Figure D: Example Networks

(A) English IPA

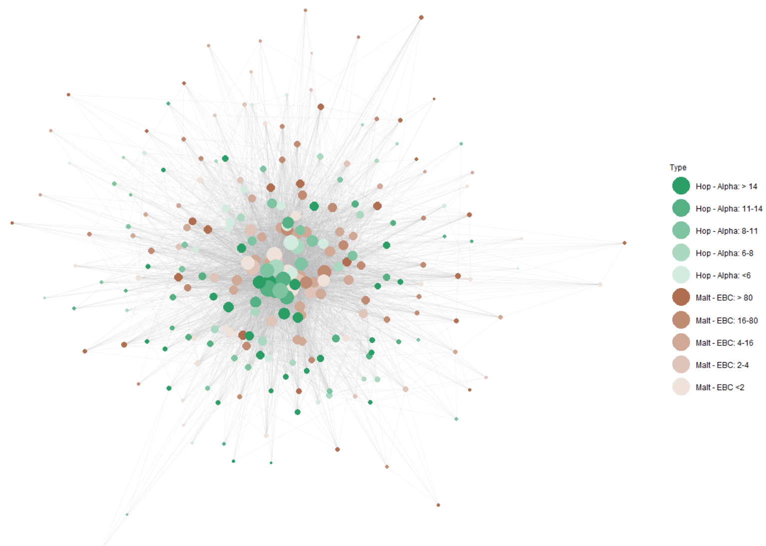

Nodes	**Density**	**Diameter**	**Clustering**	**Avg. Degree**
263	0.19	1.16	0.45	49

(B) Munich Helles

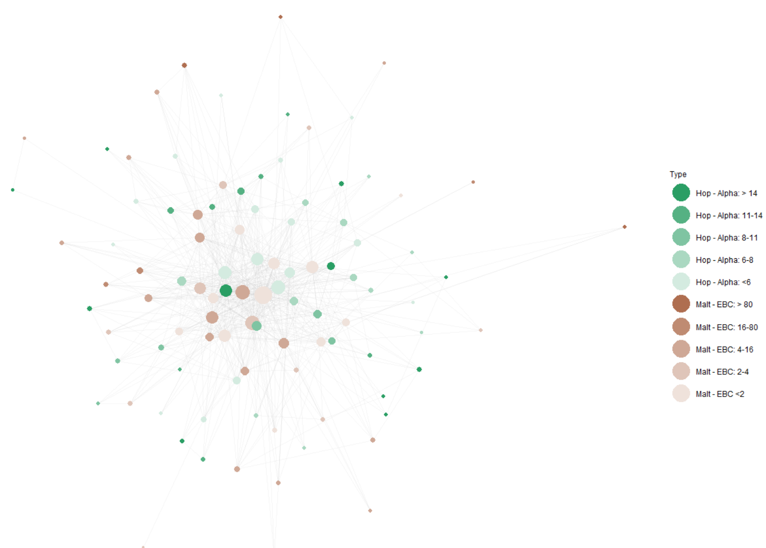

Nodes	**Density**	**Diameter**	**Clustering**	**Avg. Degree**
98	0.16	1.85	0.42	15.39

Figure E: Network Correlations

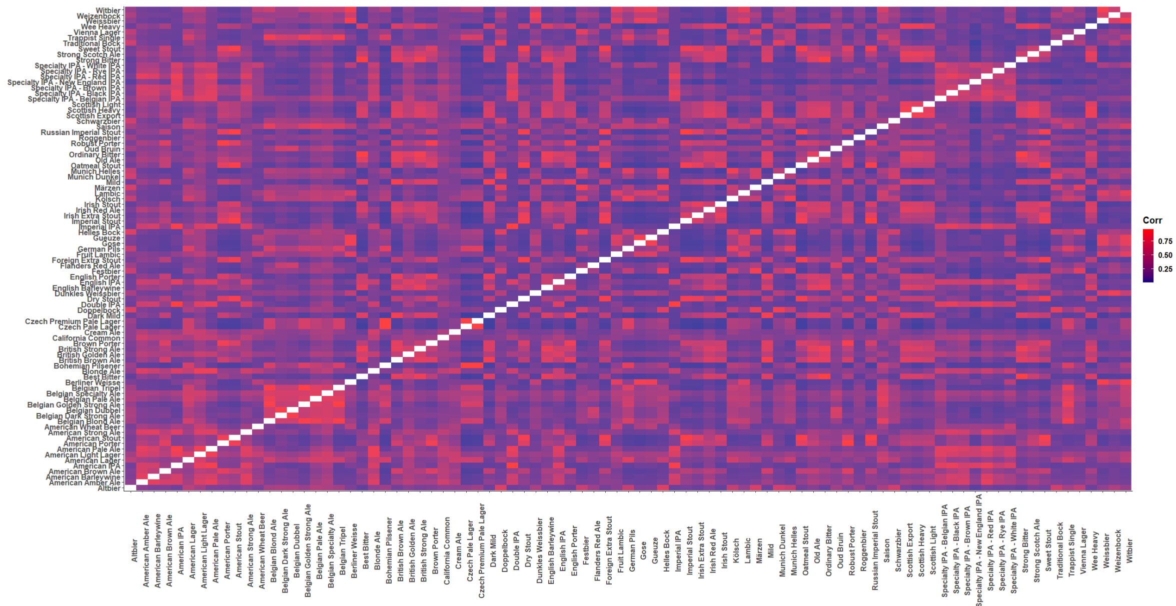

References

Abbasi, A., Altmann, J., and Hossain, L. (2011). Identifying the Effects of Co-authorship Networks on the Performance of Scholars: A correlation and regression analysis of performance measures and social network analysis measures. *Journal of Informetrics*, 5(4):594–607.

Albert, R., Jeong, H., and Barabási, A.-L. (2000). Error and Attack Tolerance of Complex Networks. *Nature*, 406(6794):378–382.

Barabási, A.-L. et al. (2016). *Network Science*. Cambridge University Press.

BarthHaas (2018). Hop Harvest Guide 2018.

BJCP (2015). Beer Judge Certification Program 2015 Style Guidelines.

Bonacich, P. (1972). Factoring and Weighting Approaches to Status Scores and Clique Identification. *Journal of Mathematical Sociology*, 2(1):113–120.

Boschma, R. (2017). Relatedness as Driver of Regional Diversification: A Research Agenda. *Regional Studies*, 51(3):351–364.

Buarque, B. S., Davies, R. B., Hynes, R. M., and Kogler, D. F. (2020). OK Computer: The Creation and Integration of AI in Europe. *Cambridge Journal of Regions, Economy and Society*, 13(1):175–192.

Cabras, I. and Bamforth, C. (2016). From reviving tradition to fostering innovation and changing marketing: The evolution of micro-brewing in the UK and US, 1980–2012. *Business History*, 58(5):625–646.

Clark, G. L., Feldman, M. P., Gertler, M. S., and Williams, K. (2003). *The Oxford Handbook of Economic Geography*. Oxford University Press.

Content, J. and Frenken, K. (2016). Related Variety and Economic Development: a Literature Review. *European Planning Studies*, 24(12):2097–2112.

Dennett, A. and Page, S. (2017). The Geography of London's Recent Beer Brewing Revolution. *The Geographical Journal*, 183(4):440–454.

Elzinga, K. G., Tremblay, C. H., and Tremblay, V. J. (2015). Craft Beer in the United States: History, Numbers, and Geography. *Journal of Wine Economics*, 10(3):242.

Feldman, M. P., Kogler, D. F., and Rigby, D. L. (2015). rKnowledge: The Spatial Diffusion and Adoption of rDNA Methods. *Regional Studies*, 49(5):798–817.

Flack, W. (1997). American Microbreweries and Neolocalism: "Ale-ing" for a Sense of Place. *Journal of Cultural Geography*, 16(2):37–53.

Frenken, K., Van Oort, F., and Verburg, T. (2007). Related Variety, Unrelated Variety and Regional Economic Growth. *Regional Studies*, 41(5):685–697.

Gatrell, J., Reid, N., and Steiger, T. L. (2018). Branding Spaces: Place, Region, Sustainability and the American Craft Beer Industry. *Applied Geography*, 90:360–370.

Healey, J. (2016). *The Hops List: 265 Beer Hops Varieties from Around the World*. Hoplist.com.

Kamada, T., Kawai, S., et al. (1989). An Algorithm for Drawing General Undirected Graphs. *Information Processing Letters*, 31(1):7–15.

Kind, C. and Kaiser, T. (2020). Heat, Hops, Hallertau: Exploring Implications of Climate Change for the German Beer Sector. In *The Geography of Beer*, pages 103–111. Springer.

Knudson, W., Sirrine, J. R., and Mann, J. T. (2020). The Geographic Dispersion of Hop Production in the United States: Back to the Future? In *The Geography of Beer*, pages 113–120. Springer.

Kogler, D. F. (2016). Evolutionary Economic Geography: Theoretical and Empirical Progress.

Kogler, D. F., Essletzbichler, J., and Rigby, D. L. (2017). The Evolution of Specialization in the EU15 Knowledge Space. *Journal of Economic Geography*, 17(2):345–373.

Kogler, D. F., Rigby, D. L., and Tucker, I. (2013). Mapping Knowledge Space and Technological Relatedness in US Cities. *European Planning Studies*, 21(9):1374–1391.

Li, X., Liu, Y., Jiang, Y., and Liu, X. (2016). Identifying Social Influence in Complex Networks: A Novel Conductance Eigenvector Centrality Model. *Neurocomputing*, 210:141–154.

Miguelez, E. and Moreno, R. (2018). Relatedness, External Linkages and Regional Innovation in Europe. *Regional Studies*, 52(5):688–701.

Mittag, R. (2014). Geographic Appellations of Beer. In *The Geography of Beer*, pages 67–74. Springer.

Molloy, M. and Reed, B. (1995). A Critical Point for Random Graphs with a Given Degree Sequence. *Random Structures & Algorithms*, 6(2-3):161–180.

Newman, M. E. (2005). Power Laws, Pareto Distributions and Zipf's law. *Contemporary Physics*, 46(5):323–351.

Nooteboom, B. (2000). Learning by Interaction: Absorptive Capacity, Cognitive Distance and Governance. *Journal of Management and Governance*, 4(1-2):69–92.

Parand, F.-A., Rahimi, H., and Gorzin, M. (2016). Combining Fuzzy Logic and Eigenvector Centrality Measure in Social Network Analysis. *Physica A: Statistical Mechanics and its Applications*, 459:24–31.

Rigby, D. L. (2015). Technological Relatedness and Knowledge Space: Entry and Exit of US Cities from Patent Classes. *Regional Studies*, 49(11):1922–1937.

Rocchetta, S. and Mina, A. (2019). Technological Coherence and the Adaptive Resilience of Regional Economies. *Regional Studies*, 53(10):1421–1434.

Ruhnau, B. (2000). Eigenvector-centrality a Node-centrality? *Social Networks*, 22(4):357–365.

Sewell, S. L. (2014). The Spatial Diffusion of Beer from its Sumerian Origins to Today. In *The Geography of Beer*, pages 23–29. Springer.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27(3):379–423.

Toth, G., Elekes, Z., Whittle, A., Lee, C., Kogler, D. F., et al. (2020). Technology Network Structure Conditions the Economic Resilience of Regions. Technical report, Utrecht University, Department of Human Geography and Spatial Planning.

Wasserman, S., Faust, K., et al. (1994). *Social Network Analysis: Methods and Applications*, volume 8. Cambridge University Press.

Whittle, A. and Kogler, D. F. (2020). Related to What? Reviewing the Literature on Technological Relatedness: Where We are Now and Where Can We go? *Papers in Regional Science*, 99(1):97–113.

Wojtyra, B., Kossowski, T. M., Březinová, M., Savov, R., and Lančarič, D. (2020). Geography of Craft Breweries in Central Europe: Location Factors and the Spatial Dependence Effect. *Applied Geography*, 124:102325.

Yool, S. and Comrie, A. (2014). A Taste of Place: Environmental Geographies of the Classic Beer Styles. In *The Geography of Beer*, pages 99–108. Springer.

UCD CENTRE FOR ECONOMIC RESEARCH – RECENT WORKING PAPERS

WP20/07 Anna Bayona and Oana Peia: 'Financial Contagion and the Wealth Effect: An Experimental Study' March 2020
WP20/08 Orla Doyle: 'Can Early Intervention have a Sustained Effect on Human Capital?' April 2020
WP20/09 Supriya Kapoor and Oana Peia: 'The Impact of Quantitative Easing on Liquidity Creation' April 2020
WP20/10 Judith M Delaney and Paul J Devereux: 'Choosing Differently? College Application Behaviour and the Persistence of Educational Advantage' May 2020
WP20/11 Jonathan Briody: 'Parental Unemployment During the Great Recession and Childhood Adiposity' May 2020
WP20/12 Tensay Meles, Lisa Ryan, and Joe Wheatley: 'COVID-19 and EU Climate Targets: Going Further with Less?' May 2020
WP20/13 Xidong Guo and Sarah Parlane: 'Addressing Private Practice in Public Hospitals' May 2020
WP20/14 Patrick Massey and Moore McDowell: 'EU Competition Law: An Unaffordable Luxury in Times of Crisis?' May 2020
WP20/15 Judith M Delaney and Paul J Devereux: 'How Gender and Prior Disadvantage predict performance in College' May 2020
WP20/16 Morgan Kelly, Joel Mokyr and Cormac Ó Gráda: 'The Mechanics of the Industrial Revolution' June 2020
WP20/17 Morgan Kelly and Cormac Ó Gráda: 'Connecting the Scientific and Industrial Revolutions: The Role of Practical Mathematics' June 2020
WP20/18 Ronald B Davies, Dieter F Kogler and Ryan Hynes: 'Patent Boxes and the Success Rate of Applications' June 2020
WP20/19 Ronald B Davies, Yutao Han, Kate Hynes, and Yong Wang: 'Competition in Taxes and IPR' June 2020
WP20/20 David Madden: 'The Socioeconomic Gradient of Cognitive Test Scores: Evidence from Two Cohorts of Irish Children' June 2020
WP20/21 Deirdre Coy and Orla Doyle: 'Should Early Health Investments Work? Evidence from an RCT of a Home Visiting Programme' July 2020
WP20/22 Cormac Ó Gráda: 'On Plague and Ebola in a Time of COVID-19' August 2020
WP20/23 Morgan Kelly: 'Understanding Persistence' September 2020
WP20/24 Jeff Concannon and Benjamin Elsner: 'Immigration and Redistribution' September 2020
WP20/25 Diego Zambiasi: 'Drugs on the Web, Crime in the Streets - The Impact of Dark Web Marketplaces on Street Crime' September 2020
WP20/26 Paul Hayes, Séin Healy, Tensay Meles, Robert Mooney, Sanghamitra Mukherjee, Lisa Ryan and Lindsay Sharpe: 'Attitudes to Renewable Energy Technologies: Driving Change in Early Adopter Markets' September 2020
WP20/27 Sanghamitra Mukherjee: 'Boosting Renewable Energy Technology Uptake in Ireland: A Machine Learning Approach' September 2020
WP20/28 Ivan Petrov and Lisa Ryan: 'The Landlord-Tenant Problem and Energy Efficiency in the Residential Rental Market' October 2020
WP20/29 Neil Cummins and Cormac Ó Gráda: 'On the Structure of Wealth-holding in Pre-Famine Ireland' November 2020
WP20/30 Tensay Meles and Lisa Ryan: 'Adoption of Renewable Home Heating Systems: An Agent-Based Model of Heat Pump Systems in Ireland' December 2020

UCD Centre for Economic Research
Email economics@ucd.ie