NHH

INSTITUTT FOR SAMFUNNSØKONOMI

DEPARTMENT OF ECONOMICS

SAM 27 2018

ISSN: 0804-6824

December 2018

Discussion paper

Feeding the people: grain yields and agricultural expansion in Qing China

BY

Liam Brunt AND **Antonio Fidalgo**

Norges Handelshøyskole

NORWEGIAN SCHOOL OF ECONOMICS

Feeding the people: grain yields and agricultural expansion in Qing China

Liam Brunt[1] & Antonio Fidalgo[2]

Abstract

We use modern econometric methods to analyze a recently-released sample of 3 000 Chinese grain yields. We find significant variation across provinces and persistent increases in yields over time – albeit slow compared to Europe and the New World. Growth rates for rice (the primary southern crop) and dry land crops (the primary northern crops) were similar. We show that provinces were more extensively farmed when yields and population pressure were high, and that extending production put downward pressure on yields. Overall, Chinese farmers avoided the problem of agricultural involution by efficiently boosting output at the extensive margin, not the intensive margin.

JEL codes: N55, O13, O47.
Keywords: agricultural involution, productivity, growth.

---

[1] NHH – Norwegian School of Economics, Department of economics, Helleveien 30, 5045 Bergen, Norway. Corresponding author. Contact: liam.brunt@nhh.no.
[2] Fresenius University of Applied Sciences.

1. Introduction. Over the last twenty years, there has been a concerted effort to improve estimates of Chinese output and productivity in order to make more accurate and meaningful comparisons with Europe (in general) and England (in particular). This is part of the process of understanding why industrialization started in England in the eighteenth century – and then spread to Europe in the nineteenth century – while China seemed to become less economically developed over the same period. This era corresponds to the reign of the Qing emperors, starting with Shunzi in 1644 and ending with the overthrow of Xuantong in 1911, and the consequent transition to a republic in 1912. The key to any economic comparison of China and Europe must lie in agriculture because 90 per cent of the Chinese population was engaged in agricultural production (Shi, 2017, p. 170, table 5.8). Hence researchers from Li (1998) through to Broadberry, Guan and Li (2017) have focused considerable attention on calculating and understanding Chinese crop yields: they are a crucial part of the puzzle. In this paper, we push the analysis further in order to present cleaner yield estimates. We then show how yields influenced – and were influenced by – the pattern of agricultural production.

Shi (2017) has done a great service to the field by producing a large, new data set of grain yields. Previous estimates were based on a sample of around 1 000 yields – which is not many, for an analysis spanning a huge country and 250 years of history! Shi carefully re-examined all the previous data and expanded the sample to around 3 000 observations, for which he presents averages in his book. However, simple averages are not very informative because there are serious sample selection issues. We use modern econometric methods to correct for these in order allow the data to speak more clearly. Our approach is very similar to the one adopted by Clark (2002) in producing a consistent rent series for English agriculture, based on a scattered sample of 32 000 rent observations spanning the period 1500 to 1912.

Shi's conclusions are very modest and rather downbeat. With regard to rice production – by far the most important component of output – he examines both the cross sectional and time series yield variation. Looking across provinces, Shi suggests (2017, p. 66) that: "generally speaking there is not much difference between the averages, the exceptions being only Jiangsu and Fujian." And looking over time he argues (2017, p. 68) that: "in most provinces averages in the early Qing are slightly lower, while averages in the Jiaqing-Daoguang (mid-Qing) are slightly higher. The comparison only has reference value: no any [sic] definite conclusions can be drawn from it alone." By contrast, we show that there was significant variation across provinces: almost every province had a unique and statistically significant time trend (not just Jiangsu and Fujian). We use these trends to generate a panel of provincial yield estimates for each Imperial reign. We then combine these estimates with data on acreage and population to examine the causes of the cross sectional variation and time series change in land productivity. We show that agriculture expanded at the extensive margin, rather than the intensive margin, and thereby avoided the classic trap of agricultural involution (Geertz, 1963).

In section 2 we discuss the data, which are quite different from the type of material that we typically handle in agricultural history and not at all straightforward. In section 3 we outline and implement our econometric analysis to produce new time-specific provincial yield estimates. In section 4 we combine the yield estimates with data on population and cultivated area to explain changes in the pattern of production. In section 5 we conclude.

2. The data. The yield data that we typically find in economic history come either from some official source (government statistics or institutional records, such as manorial rolls) or from some private account (farmers' corn books, probate inventories). They are generally "hard data" – a precise record of an output or transaction that occurred at a particular time in a particular place. For the most part, these types of records seem either to have never been generated in China or have been lost. Instead, economic historians are much more dependent on "soft data", such as reports from local gazetteers – which form the bulk of the material in this data set (Shi, 2017, p. 5). There are multiple problems with these types of data – compared to the "hard data" – which reduce their accuracy and bring into question how they should be interpreted. Some key problems are as follows.

First, many data points are highly uncertain owing to the fact that the yield must often be inferred from a related piece information (and the inference process rests on an assumption that could be very inappropriate). For example, rent books are used to estimate yields on the assumption that the rent is set at half the yield (Shi, 2017, p. 61). That is, farmers were sharecropping and the rent was recorded as a certain number of units of grain (*dan*) per unit of land (*mu*). But Shi notes that the rent was not always 50 per cent of output: it could be 40 or 30 per cent. This imparts a lot of potential measurement error into the observation: a rent of 0.5 dan could imply a yield of 1 mu (for a 50 per cent sharecrop) or 1.7 mu (for a 30 per cent share crop). As another example, the output of a piece of land may be known but not its area. However, the source may also record the amount of seed sown. Given a "typical" sowing rate, Shi can infer the area and therefore the yield per acre (Shi, 2017, p. 62). But, again, this assumption about seeding rates can induce a very substantial measurement error. The dispersion in the data suggests that measurement error is extreme for some observations, at least. The maximum rice yield is 20.00, whilst the minimum is 0.20 (and the mean is 3.02). The maximum dry land crop yield is 4.36, whilst the minimum is 0.02 (and the mean is 0.71). It is hard to see how it can have been economic to cultivate a dry land crop with a *yield* of 0.02 dan per mu (0.37 bushels per acre).[3] For comparison, farmers in England around 1800 would sow 2 bushels of wheat per acre and reap a yield of 20; and in medieval times they would sow 2 bushels and reap perhaps 10 (Titow, 1972). The latter yield is close to the Chinese average for dry land crops (0.7 dan per mu converts to 13 bushels per acre). So the average for China seems plausible but the extreme values must surely represent measurement error. We consider how to deal with this in the next section.

Even if all the observations were accurate (i.e. no measurement error) then there is still the problem of representativeness. In particular, it is generally impossible to know what the data in gazetteers represent due to our limited understanding of the situation at the time they were written. When the author wrote the gazette entry, his audience would have understood the importance of what was written because they knew the background. For example, if you were to read that in 2018 the New England Patriots lost 5 games then you would understand that this was noteworthy for being a high number; and if you read that the Miami Dolphins lost 5 games then you would understand that this was noteworthy for being a low number. But a reader 200 years from now – who did not have access to NFL performance data for previous seasons – could not interpret the data appropriately. And so it is when we are given a yield figure in a local gazette: was the author noting this because it was normal, or particularly high or low? Another problem is that the observations are often vague. We have similar records in Europe – such as the English yield estimates reported in the Board of Agriculture's county surveys around 1800. They

[3] Following Shi, 2017, p. XXII, we take 1 dan to be 103.1 litres, which equals 2.8349 Imperial bushels; and 1 mu to be 0.1515 acres.

may state: “yield on good land is 20 bushels, but only 15 on poor land”. Since we do not know the proportions of good and poor land, how should we code these data? Just take a simple average (thus implicitly assuming that good and poor land are equally prevalent)? Or disregard them completely? Western researchers effectively follow the latter strategy, in that no one is basing their analysis on data like these any more – they would be considered just too imprecise and unreliable. Yet this type of data is what we have available for China, so we have to make the best of it.

A final point to note is that the data refer either to rice or to “dry land crops”. Pooling wheat, barley, millet and so on in a category called “dry land crops” means that we are averaging across a variety of crops that have markedly different average yields. This obviously induces substantial measurement error, since a yield report based on wheat will be a lot lower (perhaps 20-30 per cent) than a yield report based on barley. Moreover, there will likely be systematic differences across provinces in terms of the frequency with which different dry land crops appear (some areas will focus more wheat production, whilst others focus more on millet production). So differences in average yields of dry land crops do not necessarily represent differences in the productivity of land. Similarly, the crop mix may change over time and this could impart trends to the data.

Fundamentally, the decision to pool all dry land crops arises from the sparseness of data. In table 1 below we report the number of observations by province. Note that the provinces in the top part of the panel produced exclusively or mostly dry land crops, so these observations are all for dry land crops; and conversely for the lower part of the panel, where the provinces focused on rice production. There are four times as many rice observations as dry land crop observations. Obviously, we are expecting to get stronger results for our analysis of rice yields, given that we have many more observations and there is likely to be less measurement error (because we are not pooling different crop types). Bearing in mind that our analysis spans 250 years, having a total of 25 observations for a province implies only one observation per decade! Heilongjiang (with three observations) commonly drops out of the analysis and Jilin (with 14 observations) is rarely statistically significantly different to the other provinces. Note that their absence is historically unimportant because Jilin and Heilongjiang were largely uncultivated for most of the Qing era. The Qing dynasty was founded by Manchu invaders, whose home was originally in the area of Jilin and Heilongjiang (i.e. Manchuria), so the Qing emperors banned the Han Chinese from going there to farm – otherwise the provinces would have been overwhelmed by migrants from the south. The restriction was relaxed only in the late nineteenth century, which is why there was very little agricultural production – and very few grain yield observations – before that time.

Table 1. The cross sectional distribution of observations.

Dry land crop provinces	N	Rice provinces	N
Zhili	95	Anhui	300
Shandong	140	Zhejiang	110
Shanxi	72	Fujian	155
Shaanxi	136	Jiangxi	213
Henan	57	Guangdong	251
Gansu	53	Jiangsu	99
Xinjiang	24	Hunan	284
Liaoning	22	Hubei	370
Jilin	14	Guangxi	224
Heilongjiang	3	Yunnan	198
		Guizhou	93
		Sichuan	251
Dry land crop total	***627***	***Rice total***	***2531***

There is another form of measurement error that we have yet to discuss – temporal measurement error. Shi dates almost all observations by the Imperial reign, not by the year. Presumably, the gazetteers typically make statements such as: “In the reign of Shunzhi, rice yields were 1 dan per mu”. Then no more precise date is available. In our analysis, we start by using reign dummies as time fixed effects, based on table 2 below. However, some reigns were much longer than others (notably, Kangxi and Qianlong each reigned for 60 years, whereas Xuantong reigned for only two years). So if we are interested in growth rates then it obviously makes more sense to think about growth per annum than growth per reign. We therefore dated each observation at the mid-point of each reign (i.e. a Kangxi observation was dated as 1692, and so on). Choosing the mid-point minimizes the expected dating error. Generating an estimate of the number of years between observations enables us to space them in a more meaningful way and allows us to estimate annual growth rates.

In fact, a year – or a short range of years – is reported for a small minority of observations. We additionally entered these data and re-analyzed everything using more precise dating when available; it made no difference. This is for two reasons. First, there are few such observations. Second, the dates that we assign to each observation are relatively close to the true date. Suppose that we date an observation to the middle of a 30-year reign. On average, the true date of the observation will randomly differ by 7.5 years from our chosen date (i.e. the true date of the observation will lie halfway between our chosen mid-point and one or other of the terminal dates for the reign). Given that we are estimating growth rates over 250 years, an error of 7.5 years for a small number of observations will make no discernible difference to the analysis.

Table 2. The temporal distribution of observations.

Emperor	Years	Reign	Rice N	Rice N/year	Dry crop N	Dry crop N/year
Shunzhi	1644-1661	1	37	2.06	30	1.67
Kangxi	1662-1722	2	246	4.03	69	1.13
Yongzheng	1723-1735	3	196	15.08	33	2.54
Qianlong	1736-1795	4	644	10.73	200	3.33
Jiaqing	1796-1820	5	334	13.36	51	2.04
Daoguang	1821-1850	6	440	14.67	56	1.87
Xianfeng	1851-1861	7	119	10.82	3	0.27
Tongzhi	1862-1874	8	210	16.15	15	1.15
Guangxu	1875-1908	9	236	6.94	109	3.21
Xuantong	1909-1911	10	61	20.33	24	8.00
Transition to Republic	1912	11	7	7	37	37.00
Undated		12	7		0	

Notes. A handful of observations are dated as "Qing dynasty"; we tag them as period 12 but make no use of that time stamp.

There is yet another dating problem. Some observations are dated to more than one reign, bearing a legend such as "Shunzhi-Kangxi" or even "Qianlong-Daoguang" (which covers three reigns). There are (at least) two ways to interpret such an entry. The first interpretation is that it is unclear to which period the observation applies. The second interpretation is that it applies to all the periods (i.e. it is implicitly an observation for each of the mentioned reigns). We coded the data both ways. Observations referencing multiple reigns were initially allocated to the reign covering the mid-point of the date range. We then created duplicate entries and allocated the observation also to the other reign(s). We tagged these duplicates so that we can run the analysis both with and without them. In total, we ended up with 228 duplicates (out of a total 3 164 entries in the data set).

It is worth noting that we are estimating rice yields in 12 provinces. So even the best documented reign, Daoguang, offers only around one yield observation per province per year. And this is for provinces that are commonly larger than European nation states! Since the data are stretched so thin, we need to treat them very carefully in order to control for measurement error and sample selection bias as best we can; taking simple averages is likely to be very misleading.

Finally, we need to consider how to match the yield data to other data (such as population and acreage) in a way that is spatially and temporally coherent. First, for the purposes of matching data, we equate Zhili (the province surrounding Beijing) with Hebei. This is not strictly accurate because borders have been redrawn (multiple times) and Beijing and Tianjin have both been removed from Hebei (redesignated as autonomous municipal authorities). However, we are not in a position to be able to adjust for these border changes; we do not believe that this significantly affects our results. Second, we take the provincial population estimates for dates when they are available (Perkins, 1969, presents data for 1749, 1776, 1819, 1851, 1873, 1893 and 1913) and we allocate them to the appropriate reign (so 1776 goes to Qianlong and 1819 goes to Jiaqing and so on). We make no attempt to adjust the population estimates for inter-censal growth (such as using the 1776 data to generate an estimate for 1766 – which is the mid-point of Qianlong's reign). Given the rough and ready nature of the data, we feel that this would just be a false attempt at precision. Provincial populations for reigns without a census were estimated by interpolation or extrapolation using the growth rate, since no alternative was available to us. Third, we would like some measure of land capacity utilization. Chinese provinces vary massively in size, with the largest (Xinjiang, with 1.6 million km^2) being 16 times the smallest (Zhejiang, with 0.1 million km^2). For comparison, France has 0.6 million km^2 and Belgium has 0.03 million km^2. However, in the largest provinces, much of the land is desert (such as the Taklamakan in Xinjiang and the Gobi in Mongolia). So surface area is no guide to agricultural potential. We therefore take the agriculture census of 1952 as a benchmark. By that time, the civil war was over and the new Communist government was making a push to maximize output. China was still predominantly agricultural. Hence it seems reasonable to suppose that any land that was economic to farm (and maybe some that was uneconomic) was in production in 1952. Around 60 per cent land in prime agricultural provinces (Shandong, Jiangsu) was in production in 1952, but less than one per cent of Xinjiang. This demonstrates the importance of having a metric of agricultural potential that is independent of simple province size. Hence we take the agricultural census of 1952 as our estimate of cultivable land in each province.

3. Estimating provincial grain yields. Given the measurement error problems, the yield data are not normally distributed: there is too much mass in the tails and the distribution is right-skewed. (A right skew is almost inevitable because the data have a zero lower bound but no upper bound.) The obvious solution is to take natural logarithms because this crushes the outliers – especially the right tail – and makes the distribution much closer to normal. This reduces the influence of points of high leverage (which are exactly the ones that are most likely to be subject to measurement error, and which can otherwise seriously bias our econometric estimates). Taking logs also has the advantage that we can then interpret the regression coefficients as percentage changes or growth rates. This is the approach that we take. An alternative would be to "clean" the data set somehow. The problem is that we have no independent information on which data points need "cleaning". So then we end up inferring that certain observations are erroneous simply because they are outliers and we arbitrarily choose a rule to cut them off (i.e. we truncate the distribution). We find this to be an unattractive way to proceed and thus avoid it if possible. It turns out that taking logs works very well in our setup.

There is one further point to note. Although the yield data suffer from a lot of measurement error, we will be using them as our dependent variable. Measurement error in the dependent variable degrades the fit of the regression (reduces the r-squared) but leaves the coefficients on the explanatory variables unbiased. This is an important point. Obviously, it gives us more confidence in the accuracy of our model (since it is inherently unbiased and thus we do not need to use any sophisticated econometrics to "fix" it). But, more than that, at the end of this process we are going to end up with two panels of grain yields – one for rice, one for dry land crops – because we are going to generate estimates for each province in each reign using the two new models (one for rice, one for dry land crops). Crucially, since our panel of yields will be based on our econometric model – which is not biased by the measurement error – our panel of yields will not be contaminated by the measurement error in the underlying data. By contrast, taking simple averages of underlying data that contain a lot of measurement error *would* generate biased results unless the errors are symmetrically distributed (which is very unlikely, given that they are right-skewed). So this is a case in which yield estimates based on a model are plausibly more accurate than estimates based directly on the raw data.

Our first objective is simply to use econometrics to describe the data accurately. What does this mean? We know from tables 1 and 2 that some provinces, and some time periods, have far more observations than others. If yields are growing then the distribution of observations across space and time will affect the observed average yield. If all the observations for Province A are drawn from 1700, and all the observations for Province B are drawn from 1900, then it is almost certain that the average yield in Province B will be higher. So we need to control for temporal variation – which itself may vary across provinces – in order to make clean comparisons of average yields across space and time. We are therefore going to start with a regression such as:

$$\ln Yield_{it} = a + b.Province_i + c.Period_t + d.Province*Period_{it} + e_{it} \quad (1)$$

All our analysis will be done separately for rice and dry land crops. Since data on the two crops do not overlap spatially – i.e. we have rice yields for some provinces and dry land crop yields for the others – there is really nothing to be gained from pooling them in the same regression. Note that this is not a panel regression. The data do not yet have a regular structure, with the same number of cross sectional units at each benchmark date. This is just a pooled OLS with 2 531 yield observations on the left hand side (627 for dry land crops). It is more like a hedonic regression. Once we have estimated this regression, we can use the coefficients to create a panel of yields (i.e. an estimated average yield for each province-period) and this will be used in the next section of the paper.

We used several approaches to modeling time, as reported in table 3 below. An obvious approach is to have a dummy for each reign (i.e. use time fixed effects), since that is how the observations were originally recorded. It turns out that most of the fixed effects are not statistically significant; and the coefficients on the fixed effects are generally small near the beginning and large near the end (i.e. yields were rising secularly). It turns out that replacing time dummies with a continuous variable is a superior solution (it is more efficient to estimate one continuous variable than ten dummy variables, and the pace of change was sufficiently regular that this does little violence in describing the data). The continuous variable can either be a simple count variable (i.e. just ticking up from period 1 to period 10) or a time variable (i.e. using the year at the mid-point of each reign, which we previously assigned to each observation). The equivalence of these approaches can be seen in table 3 by comparing model 1 (time dummies) to model 2 (a count variable) and model 3 (year). Using year obviously has the advantage that the coefficient is naturally interpreted as an annual growth rate, which turns out to be a mere 0.1 per cent per annum. We estimated all these models using a general-to-specific approach (i.e. we added all the province dummies and time dummies and interaction terms and then progressively eliminated the least significant ones until we were left only with variables that were statistically significant).

Table 3. Describing the cross sectional and temporal pattern of rice yields with OLS.

lnYield	Model 1	Model 2	Model 3	Model 4
Period2	0.0845* (0.0369)			
Period5	0.0728* (0.0325)			
Period6	0.1453** (0.0299)			
Period8	0.2521** (0.0393)			
Period9	0.1792** (0.0378)			
Period		0.0245** 0.0050		
Year			0.0009** (0.0002)	0.0015** (0.0002)
Jiangsu	-0.2441** (0.0606)	-0.2222** (0.0607)	-0.2194** (0.0608)	8.1961** (1.2682)
Anhui	-0.3128**	-0.3000**	-0.3044**	

	(0.0441)	(0.0439)	(0.0440)	
Zhejiang	0.0113 (0.0584)	0.0350 (0.0586)	0.0357 (0.0587)	2.6787* (1.4030)
Fujian	0.0885 (0.5219)	0.1120* (0.0523)	0.1157* (0.0524)	
Jiangxi	-0.2240** (0.0487)	-0.1942** (0.0484)	-0.1943** (0.0484)	
Guangxi	-0.5127** (0.0474)	-0.5113** (0.0479)	-0.5086** (0.0479)	
Hubei	-0.5533** (0.0419)	-0.5604** (0.0422)	-0.5602** (0.0422)	
Hunan	-0.1493** (0.0443)	-0.1539** (0.0445)	-0.1535** (0.0446)	
Sichuan	-0.1883** (0.0459)	-0.1988** (0.0467)	-0.1960** (0.0468)	
Yunnan	-0.4733** (0.0488)	-0.4687** (0.0492)	-0.4623** (0.0494)	3.8130** (1.5862)
Guizhou	-0.5103** 0.0634	-0.5022** 0.0622	-0.4997** 0.0633	
Jiangsu*year				-0.0047** (0.0007)
Anhui*year				-0.0002** (0.0000)
Zhejiang*year				-0.0015* (0.0008)
Fujian*year				0.00007** (0.00003)
Jiangxi*year				-0.0001** (0.00003)
Guangxi*year				-0.0003** (0.00003)
Hubei*year				-0.0003** (0.00002)
Hunan*year				-0.0001** (0.00002)
Sichuan*year				-0.0001** (0.00003)
Yunnan*year				-0.0024** (0.0009)
Guizhou*year				-0.0003** (0.00003)
Constant	1.1599** (0.0339)	1.1053** (0.0398)	-0.3305 (0.3184)	-1.4336** (0.3467)
r^2	0.16	0.15	0.15	0.17
Adjusted-r^2	0.16	0.15	0.15	0.16
N	2537	2537	2537	2537

Notes. Period is a count variable, running from 1 to 11. ** denotes statistical significance at the 1 per cent level; * denotes statistical significance at the 5 per cent level. Standard errors are in brackets. The base (missing) province is Guangdong. The Zhejiang dummy is retained in models 1 to 3 for completeness; it makes no difference to the estimated equations if it is dropped.

It is reassuring to find that the results are robust to different formulations of the time trend. We then extended the analysis by re-estimating the equations with provincial time trends, as shown in model 4. There are several points to note here. First, all the provincial time trends are negative – but this does not generally mean that yields were falling. It simply means that they were growing more slowly than the missing province (Guangdong). So Guizhou yields were growing at 0.0015-0.0003=0.0012 (i.e. 0.12 per cent per annum). Second, note that most of the province dummies have dropped out. Why? Because the differences in the average provincial yields were not due to a static difference between them; rather, they were due to the fact that they were growing at different rates. Suppose that the yields in Guangdong and Fujian started at the same level but Fujian yields ended 23.14 per cent higher. Then the average level of Fujian yields would be 11.57 per cent higher over the period. Then – as in model 3 – the coefficient on the Fujian dummy will be 0.1175. If we let the provincial yields differ *only* in terms of their level – as in models 1 to 3 – then the faster growth of Fujian yields must be absorbed by the Fujian province dummy. But this is not the best way to model difference between Fujian and Guangdong. It turns out that there is no significant difference between them at the beginning but they follow different trends; hence the Fujian dummy becomes statistically insignificant when we allow the Fujian trend to be greater than the Guangdong trend. The fact that virtually all the provincial dummies (fixed effects) become insignificant – but all the provincial time trends are significant – shows that it is important not to just look at static averages. Contrary to Shi's inference from the raw data, all the provinces had significantly different average yields (as shown in models 1 to 3) caused by different trends (as shown in model 4).

Let us now turn to dry land crop yields, where we repeat the process that we used to model rice yields and report the output in table 4 below. The results for dry land crops are similar to rice but less tidy. For example, when we move from time fixed effects (model 1) to a period trend (model 2) then model fit deteriorates – the r-squared falls and the coefficient (and hence statistical significance) falls on some of the province dummies. A time trend captured by year (model 3) is statistically significant only at the 15 per cent level (i.e. significantly worse than the period trend). Introducing provincial time trends substantially improves the fit (model 4) and makes the time trend (year) highly statistically significant (model 5). The latter is our preferred model; the overall structure and fit and are similar to our preferred model of rice yields, and it is straightforward to interpret. It makes little difference whether we use model 4 or 5 to construct a panel of predicted province-period grain yields (the results are correlated around 0.85). Note that Gansu has an incredibly large province fixed effect – bearing in mind that the yields are in natural logarithms – and this is a point that we examine in more detail later.

Table 4. Describing the spatial and temporal pattern of dry land crop yields with OLS.

lnYield	Model 1	Model 2	Model 3	Model 4	Model 5
Period 3	0.3320* (0.1661)				
Period 4	-0.2323** (0.0839)			-0.2307** (0.0873)	
Period 6	-0.3447** (0.1250)			-0.2784* (0.1269)	
Period 8	0.5585** (0.2227)			0.6034** (0.2229)	
Period 9				0.2305* (0.1033)	
Period 10	0.6182** (0.2185)			0.5676** (0.1982)	
Period		0.0265* (0.0129)			
Year			0.0007 (0.0005)		0.0016** (0.0005)
Zhili	-0.2215* (0.1097)	-0.2287* (0.1149)	-0.2105 (0.1150)		
Henan	0.1047 (0.1316)	0.0592 (0.1327)	0.0690 (0.1329)	6.0040* (2.7979)	5.6059* (2.9124)
Shanxi	-0.5422** (0.1228)	-0.5710** (0.1216)	-0.5740** (0.1218)		
Shaanxi	0.2857** (0.1029)	0.2218* (0.1004)	0.2222* (0.1007)		
Gansu	-0.5847** (0.1429)	-0.4929** (0.1356)	-0.4979** (0.1359)	14.1015** (3.3034)	16.0606** (3.4045)
Xinjiang	0.6906** (0.1949)	0.5366** (0.1867)	0.5251** (0.1869)		
Liaoning	-0.5803** (0.2182)	-0.4003* (0.2004)	-0.3636 (0.1997)		
Jilin	-0.3597 (0.2557)	-0.1417 (0.2391)	-0.1244 (0.2396)		
Heilongjiang	1.0069** (0.4939)	0.6875* (0.4948)	0.6800 (0.4961)		
Zhili*year				-0.0001* (0.0001)	-0.0001* (0.0001)
Henan*year				-0.0032* (0.0015)	-0.0031* (0.0016)
Shanxi*year				-0.0003** (0.0001)	-0.0003** (0.0001)
Shaanxi*year				0.0002** (0.0001)	0.0001* (0.0001)
Gansu*year				-0.0082** (0.0019)	-0.0094** (0.0019)
Xinjiang*year				0.0004** (0.0001)	0.0003** (0.0001)
Liaoning*year				-0.0003** (0.0001)	-0.0002* (0.0001)
Heilongjiang*year				0.0005* (0.0003)	0.0003 (0.0003)
Constant	-0.5759** (0.0705)	-0.7601** (0.0932)	-1.877* (0.8586)	-0.6426** (0.0723)	-3.5454** (0.9254)
r^2	0.17	0.12	0.11	0.19	0.15
Adjusted-r^2	0.15	0.10	0.10	0.17	0.13
N	627	627	627	627	627

Notes. Period is a count variable, running from 1 to 11. ** denotes statistical significance at the 1 per cent level; * denotes statistical significance at the 5 per cent level. Standard errors are in brackets.

Are the yield growth rates that we have estimated plausible? Was growth "fast" or "slow"? At this point, it is useful to compare the Chinese results to the experience of other countries, as in table 5 below. It is surprisingly difficult to find yield data for European countries around 1700; we present all the data that we have been able to find. On average, yield growth in western countries was six to twelve times higher than yield growth in China between 1700 and 1914. If we exclude extreme observations in China (i.e. Jiangsu and Gansu) then the average growth rate of yields rises to around 0.12 per cent per annum for both rice and dry land crops. If we exclude extreme observations in the western nations (Prussia) then the average rate of growth falls to 0.28 per cent per annum, which is still more than twice the rate of Chinese growth. (Moreover, there is no good reason to exclude Prussia: the numbers for rye – the main field crop – are sound and wheat yields rose similarly fast, suggesting that this was a real phenomenon. In fact, Prussian success was based on adopting a very capital-intensive form of cultivation by 1914, using large amounts of artificial fertilizer; see Grant, 2005, for a detailed discussion.) Overall, we can say that the growth in Chinese grain yields was very slow by international standards.

Table 5. Growth rates of regional crop yields.

Rice:	**Growth (% p.a.)**	*Dry land crops:*	**Growth (% p.a.)**	*Wheat:*	**Growth (% p.a.)**
Guangdong	0.15	Shandong	0.16	England	0.32
Jiangsu	-0.32	Zhili	0.15	France	0.06
Anhui	0.13	Shanxi	0.13	Ireland	0.37
Zhejiang	0.00	Shaanxi	0.15	Prussia (rye)	1.07
Fujian	0.14	Gansu	-0.77	Spain	0.27
Jiangxi	0.14	Xinjiang	0.19	USA	0.42
Guangxi	0.12	Liaoning	0.14	South Africa	0.21
Hubei	0.12	Henan	-0.13		
Hunan	0.14	Jilin	0.15		
Sichuan	0.14	Heilongjiang	0.13		
Yunnan	-0.01				
Guizhou	0.12				
MEAN	***0.06***	***MEAN***	***0.03***	***MEAN***	***0.39***
ST DEV	***0.14***	***ST DEV***	***0.30***	***ST DEV***	***0.32***

Notes and sources. Chinese growth rates are calculated by adding the province-specific trends to the general time trend, as estimated in tables 3 and 4 above. Other growth rates are calculated from yield data appearing in the following sources. England: Overton, *Agricultural revolution*, p. 77; and British Government, *Returns of Agriculture*. France: Toutain, *Produit*; and Bennet, "Trends", p. 70. Ireland: Petty, *Political anatomy*, pp. 57-8; and Government of Ireland, *Farming*, table E. Prussia (where the main crop is rye): Hagen, *Ordinary Prussians*, pp. 212 and 314; and Grant, *Migration*, p. 228. South Africa: Van Duin and Ross, *Economy* (assuming the yield per acre – like yield per seed – was the same in 1700 as in 1833-42) and Nhemachena and Kirsten, "Historical assessment", figure 1. Spain: Bringas Gutiérrez, *Productividad*, p. 24 and Simpson, "Spanish agriculture", table 8.

4. Explaining provincial grain yields. Having estimated yield models for rice and dry land crops, we used them to predict yields in each province in each reign. That is, we transform the 3 000 individual yield observations into two panels of yields. The rice panel contains 11 cross sections (one for each reign) of 12 provinces (so 132 observations in total). The dry land crop panel contains 11 cross sections of 10 provinces (so 120 observations in total, although with a number of missing observations for Jilin and Heilongjiang). For the convenience of other researchers, we report these predicted yields in appendix 1. What can we learn from these yields? One question is whether we can explain the observed yield variation. A second question is whether these yields explain any other economic variables.

We begin by noting that it is relatively challenging to explain crop yields. Suppose that we have a standard Cob-Douglas production function for rice:

$$Q_{it} = A_{it} L_{it}^{\alpha} N_{it}^{\beta} K_{it}^{1-\alpha-\beta} \quad (2)$$

where Q is output, A is technology, L is land, N is the number of workers and K is capital. We are starting with yield, Y, so we can construct Q by multiplying it by land (=Y.L). If we then use regression analysis to estimate (2) then we have effectively put L on both sides of the equation – literally so, when we take natural logarithms – so the r-squared will be high and everything will look nice. But this does not mean that the equation genuinely explains very much of the underlying variation. It is just telling us that large provinces (which generally have larger L) have more output (larger Q), which is rather obvious (especially given the gargantuan variation in province sizes).

Suppose that we instead estimate a model of yields:

$$Y_{it} = Q_{it} L_{it}^{-1} = A_{it} L_{it}^{\alpha-1} N_{it}^{\beta} K_{it}^{1-\alpha-\beta} \quad (3)$$

The r-squared on this equation will certainly be lower (since it is not artificially inflated by adding L to both sides). But now we are asking a more interesting question: to what extent is yield a function of the intensity of labour and capital use? This relates to the issue of "agricultural involution" (Geertz, 1963). This is the idea that population pressure causes intensification of agricultural production – thereby raising output per unit area but reducing output per worker. Given that China seems to have fallen behind Europe in the Qing period (Broadberry et al., 2017), and given that the population was rising fast (Perkins, 1969), it would be interesting to know the extent to which rising wheat yields were driven by rising labor input per unit area. Note that we can also add L to the set of explanatory variables, which is effectively a test for economies of scale. We can also add the *change* in L. Why would we want to do this? A standard assumption is that farmers colonize the most productive land first in any locality. So as the area of land in production rises, it puts downward pressure on average yield (since marginal yield is below average yield). This would be consistent with a Malthusian model of population: increasing population induces the next generation of native-born – or immigrants – to bring outlying (less productive) land into production. Chen and Kung (2016) provide evidence that the Malthusian model of population offers a good description of Chinese population in this period. Consistent with their characterization, we see – in table 6 below – that increasing land area turns out to be an important explanator of changes in yields.

Table 6. Explaining the pattern of rice yields.

lnYield	Model 1 (OLS)	Model 2 (OLS)	Model 3 (IV)
Δ.Percentage of agricultural area in use	-0.0027* (0.0013)		-0.0107 (0.0063)
Δ.lnAgricultural area		-0.0987* (0.0487)	
lnPopulation/agricultural area in use	0.0292 (0.0240)	0.0262 (0.0243)	-0.0280 (0.0539)
Constant	1.2092** (0.1969)	1.1827** (0.2007)	0.7887* (0.4088)
Within group r-squared	0.07	0.08	0.06
Between group r-squared	0.40	0.22	0.05
Overal r-squared	0.09	0.07	0.03
N	121	121	110

Notes. ** denotes statistical significance at the 1 per cent level; * denotes statistical significance at the 5 per cent level. Standard errors are in brackets. In Model 3 we instrument for Δ.*Percentage of agricultural area* *in use* using lagged values and province fixed effects and time fixed effects; the instrumented variable is statistically significant at the 10 per cent level.

We can model the change in the intensity of land use in two ways. Either we can use the change in physical area (model 2), or the change in the percentage of agricultural land in use (model 1). Both give significant results but percentage of agricultural land in use gives a much better fit. This is not surprising because some provinces are much larger than others – so it takes a much larger absolute change in the area to have the same impact on yields. It is reassuring that changes in land use explain rice yields. Note that the acreage data are drawn from completely different sources to the yield data (the acreage data are provincial estimates derived from tax records), so nothing in the data construction implies that they will automatically be correlated. Unfortunately, we do not have any data to estimate the impact of capital intensity. The other variable of interest – labor intensity – is not significant, although it has the expected sign. So there is no strong evidence of involution (i.e. increasing the productivity of land by cultivating it more intensively using extra labor). Note that it would be normal to add fixed effects and time fixed effects to a panel regression, but this makes no sense here. The yield data were constructed using fixed effects and time fixed effects, so those dummies will merely extract the variation that we put into the estimates when we constructed them.

Now let us consider the extent of agricultural production. That is, how can we explain how much land is cultivated? Again, provinces that have more agricultural land will typically have a larger physical area in production, so that it not a useful metric. Instead, consider the percentage of agricultural land in production in each province. What do we expect to find? Areas that have high yields should have a high percentage of land in production. Why? Thinking about cross sectional variation, high yields are mostly determined by climate and topography. They could also be determined by regional variations in technology (such as local rice varieties). In areas with exogenously high yields, farmers will extend production, *ceteris paribus*. Notice that we are charting a relationship here between the level of yields and the level of land utilization (not the *change* in land utilization, as in our earlier analysis). We would also expect that high population density – such as the presence of many urban consumers in the province, compared to the available agricultural area – will prompt farmers to bring a higher proportion of agricultural land into production. This is exactly what we see in table 7 below.

Table 7. Explaining the percentage of agricultural land in production (rice areas).

Percentage of agricultural area in use	Model 1 (OLS)	Model 2 (IV)
lnYield	39.14** (9.82)	38.94** (11.07)
lnPopulation/total agricultural area	17.47** (1.82)	17.16** (2.16)
Constant	196.28** (21.41)	194.29** (24.44)
Within group r-squared	0.57	0.50
Between group r-squared	0.58	0.60
Overall r-squared	0.56	0.55
N	132	120

Notes. ** denotes statistical significance at the 1 per cent level; * denotes statistical significance at the 5 per cent level. Standard errors in brackets. In Model 2 we instrument for *lnYield* using lagged values, province fixed effects and provincial time trends.

A question might arise at this point: is it legitimate to regress the amount of land on yield (table 7), as well as the yield on the (change in) the amount of land (table 6). Does this not imply that endogeneity is present in both of the regressions? Not necessarily. In table 7 we are estimating a relationship in levels, whereas in table 6 we are estimating a relationship between a level and a rate of change. There are other examples in economics with a similar structure. Notably, the standard textbook macroeconomic model (ISLM) predicts a *positive, causal* impact of investment on the price level: high investment implies high aggregate demand and thus a high price level. Yet many papers also predict a *negative, causal* impact of the change in the price level (i.e. inflation) on investment: businessmen do not like uncertainty, so price level volatility discourages investment. A recent empirical paper demonstrating this result is Barro (2013), where he employs a panel data set with similar dimensions to ours to show that higher inflation reduces investment. He is obviously aware of the possibility of reverse causation and therefore uses lagged values to instrument for the (possibly) endogenous variables, such as inflation. We follow him by using lagged values as instruments, as well as province fixed effects and time fixed effects. Like him, we find no significant difference between the instrumented and non-instrumented equations (as revealed in the last columns of tables 6 and 7 above).

Pursuing the same examination of the dry land crop yields generates essentially the same results as for the rice yields. Table 8 below reveals that increasing the percentage of agricultural land in use drives down the average yield of dry land crops. At the same time, table 9 shows that high yields draw a high proportion of cultivable land into production, as does a high provincial population density (per unit of available agricultural land). The coefficients (i.e. elasticities) are very similar for rice and dry land crops, which suggests that we are quantifying real effects.

Table 8. Explaining the pattern of dry land crop yields.

lnYield	Model 1 (OLS)	Model 2 (OLS)	Model 3 (IV)
Δ.Percentage of agricultural area in use	-0.0047* (0.0022)		-0.0330** (0.0097)
Δ.lnAgricultural area		-0.0774** (0.0261)	
lnPopulation/agricultural area in use	0.0330 (0.0274)	0.0115 (0.0251)	-0.1160 (0.0625)
Constant	-0.4468** (0.1660)	-0.0557** (0.1984)	0.1281 (0.3895)
Within group r-squared	0.10	0.12	0.15
Between group r-squared	0.55	0.01	0.65
Overall r-squared	0.07	0.02	0.03
N	77	77	69

Notes. ** denotes statistical significance at the 1 per cent level; * denotes statistical significance at the 5 per cent level. Standard errors are in brackets. In Model 3 we instrument for Δ.*Percentage of agricultural area in use* using lagged values, province fixed effects and provincial time trends.

Table 9. Explaining percentage of agricultural land in production (dry crop areas).

Percentage of agricultural area in use	Model 1 (OLS)	Model 2 (IV)
lnYield	24.83** (9.94)	29.75** (11.20)
lnPopulation/total agricultural area	13.77** (1.72)	13.56** (1.94)
Constant	29.76** (12.53)	34.44** (14.12)
Within group r-squared	0.56	0.54
Between group r-squared	0.52	0.46
Overall r-squared	0.57	0.51
N	83	77

Notes. ** denotes statistical significance at the 1 per cent level; * denotes statistical significance at the 5 per cent level. Standard errors are in brackets. In Model 2 we instrument for *lnYield* using lagged values, province fixed effects and provincial time trends.

There is one caveat to the preceding analysis of dry land crops: we dropped Gansu. Yields in Gansu plummeted in the nineteenth century, as evidenced in figure 1 below, where we plot all the available yield observations contained in the raw data set. The Gansu trend is so strongly negative – in marked contrast to the mildly positive trends in virtually all the other provinces – that including Gansu makes everything statistically insignificant and sometime flips the signs in the analysis. Why is the behavior of Gansu so different? There are two unusual historical developments in Gansu that are likely related. First, the percentage of agricultural land in production in Gansu was incredibly high in the late nineteenth century. Using the 1952 acreage data as an upper bound for the area of economically cultivable land in each province, Gansu had already reached 101 per cent by 1836. Thereafter, it climbed to 112 per cent (1856), 121 per cent (1868) and 130 per cent (1892), whereupon it leveled out. By contrast, the percentage of cultivable land that was in production in the other provinces was tightly clustered around 95 per cent by 1912. With such an extensive area under cultivation in a province that is largely desert – i.e. a province in which marginal land was likely of very low quality – it is perhaps no surprise that yields should have been pushed so low in Gansu. The percentage of land in production increased by 100 per cent from 1653 to 1912 (from 31 to 131 per cent), which implies a 47 per cent fall in yields (-0.0047*100=-0.47, from table 8 above). This would reduce expected yields from 0.78 to 0.41 dan/mu – as compared to an observed reduction to 0.10 dan/mu (appendix table A2) – and so explains more than half of the observed decline.

Could this have been rational? First, note that there was an overall increase in output despite the fall in yields. A 30 per cent increase in the amount of land in production after 1836 implies a yield reduction of 14 per cent (-0.0047*30=-0.141) but this was greatly outweighed by the increase in acreage, which generated an overall increase of 16 per cent (i.e. 30-14=16). Second, note that the acreage expansion may not have been a purely economic phenomenon. The Dungan revolt broke out in 1862 in Gansu province; Zuo Zongtang was sent with an army in 1867 to crush it. He made eastern Gansu his center of operations and gradually pushed the rebels further west, into Xinjiang, where they were finally defeated in 1878. Zuo then continued to rule Gansu and Xinjiang, on behalf of the Emperor, until his death in 1885. As well as being a successful general, Zuo was a noteworthy agricultural reformer. In particular, he pressured local cultivators to cease opium production and substitute other crops instead; he carried out agricultural experiments, wrote pamphlets and worked on outreach to local farmers. Hence the presence of an intensive military-political-economic pacification campaign may explain the odd evolution of both acreage and yields in Gansu.

Figure 1. Raw yield data on Gansu dry land crops.

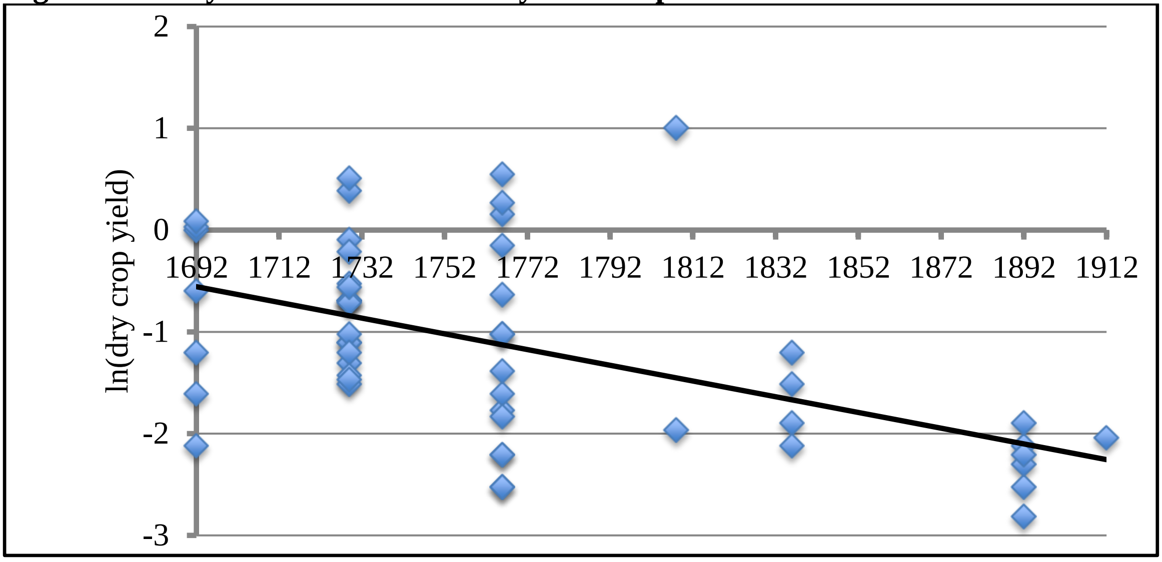

Again, all our results run counter to the agricultural involution model. In the face of rising population pressure, Chinese farmers responded during the Qing dynasty by increasing the amount of land in production (i.e. they reacted on the extensive margin, not the intensive margin). So the agricultural sector was not burdened by excessive population growth. The extensive margin increased by 40 percentage points in rice areas (that is, the average percentage of agricultural land in production increased from 54 to 94 per cent in the 250 years between the reign of Shunzi and the birth of the Republic in 1912). From table 6, we can see that this 40 percentage point increase implies a reduction in yields of 11 per cent (-0.0027*40=-0.108). So the first order effect of increasing acreage (i.e. a 40 per cent increase in output) was far greater than the second order effect of pushing down yields (i.e. an 11 per cent fall in output). Moreover, notice that the downward pressure on yields was not primarily responsible for the disparity in yield growth between China and the rest of the world. Table 5 shows that yields actually rose by 16 per cent on average (1.0006^250=1.16) over the period. Even if Chinese yields had risen by 27 per cent (=16+11 per cent) then yield growth abroad would still have been more than twice as fast as China, on average.

The analogous calculation for dry land crops shows a similar increase in the extensive margin of 45 percentage points (that is, the average percentage of agricultural land in production increased from 50 to 95 per cent over the same period – excluding Gansu and the provinces for which there are no early data, such as Heilongjiang, Jilin and Xinjiang). From table 8, we can see that this 45 percentage point increase implies a reduction in yield of 21 per cent (-0.0047*45=-0.211 per cent). So the first order effect of increasing acreage (i.e. a 45 per cent increase in output) was substantially greater than the second order effect of pushing down yields (i.e. an 21 per cent fall in output). Table 5 shows that yields actually rose by 28 per cent on average (1.001^250=1.28) for these provinces. Even if Chinese yields had risen by 49 per cent (=28+21 per cent), western yield growth would still have been twice as fast (or more).

Conclusion. Shi has performed a great service by compiling and presenting a new data set of 3 000 observations of grain yields. However, his analysis does not make best use of the material. Growth in Chinese grain yields was more widespread and more persistent than Shi infers, in both rice and dry land crops. It is nonetheless true that Chinese yield growth was only a fraction – a quarter or less – of the rate observed in western countries over the same period. Some of the differential can be ascribed to the drastic increase in Chinese acreage, which brought more marginal land into production and put downward pressure on yields. Compare the situation in China with that of the western nations. One group of western nations (such as England and France) brought very little extra land into production (about 20 per cent and 5 per cent respectively) because the available cultivable land was already largely exploited in 1700. Another group of western nations (such as South Africa and the USA) were massively increasing the amount of land in production but it was all new land (i.e. it was not marginal: there is no reason to believe that California is inherently less productive for agriculture than New England – indeed, quite the reverse). By contrast, China faced the challenge of massively expanding cultivation in regions where the best land had already been taken into production hundreds, or perhaps thousands, of years earlier.

Even controlling for the significant marginal land effect, we still need to explain half – or more – of the yield growth differential between China and the west. This comes down to capital and technology, which are distinct inputs that can also interact with one another in important ways. Prussia transformed its yield performance by adding fertilizer (which is a form of capital). But it is noteworthy that the type of fertilizer employed was one that Prussia invented itself (which was thus a new technology): it could not rely on imports of cheap South America guano, like England, so Prussia created artificial nitrogenous fertilizer. The USA invested heavily in finding – and creating – high performance seed varieties that yielded well in the harsh environment of the Great Plains. Despite the interest and energy of Zuo Zongtang in Gansu, this line of progress seems to have been largely absent from China.

One success that China can claim, in contrast to Indonesia or other Southeast Asian countries, is the avoidance of agricultural involution – that is, the intensification of production by the devotion of an ever larger workforce to a limited arable area. This prevented output per worker falling as much as it might have done in China. In fact, our results suggest that both rice and dry land crop yields were essentially invariant to labor intensity at the levels we observe in the data, so increasing the labor force would have been more an exercise in redistributive taxation than output maximization. That is to say, increasing the workforce density (for example, by splitting the farm equally between all the children in each generation, as they did in France) would have led to an agricultural pie of the same size being distributed between an ever larger number of people. So this would be purely a policy of output redistribution, rather than output expansion. Instead, it is more economically efficient (i.e. maintains output per worker) to make the younger sons leave the land completely (as in England, where they went to the cities), or have them move to new agricultural districts to expand the acreage in production (as in China).

Although we have been able to say something about the productivity of land and labor, we are missing important information about Chinese capital and technology. There was inevitably a variety of production systems in operation across a country as vast as China – not only encompassing wheat- and rice-based agriculture, but also single- and double-cropping systems, and highland versus lowland conditions. An investigation of these factors lies beyond the scope of this paper but would surely be an important further step in our understanding.

Appendix 1. Estimated grain yields in Qing China.

Table A1. Rice yields (dan/mu).

	1653	1692	1729	1766	1808	1836	1856	1868	1892	1911	1912

Guangdong	2.79	2.96	3.13	3.31	3.52	3.67	3.78	3.85	3.99	4.10	4.11
Jiangsu	4.01	3.54	3.14	2.78	2.43	2.22	2.08	2.00	1.85	1.74	1.73
Anhui	2.08	2.19	2.30	2.42	2.55	2.65	2.72	2.76	2.85	9.92	2.93
Zhejiang	3.56	3.56	3.56	3.56	3.57	3.57	3.57	3.57	3.57	3.57	3.57
Fujian	2.49	2.63	2.77	2.92	3.10	3.22	3.32	3.37	3.49	3.59	3.59
Jiangxi	2.31	2.44	2.57	2.70	2.86	2.97	3.06	3.12	3.21	3.30	3.30
Guangxi	1.71	1.80	1.88	1.96	2.06	2.13	2.18	2.22	2.28	2.33	2.33
Hubei	1.64	1.72	1.80	1.87	1.97	2.03	2.08	2.11	2.17	2.22	2.22
Hunan	2.39	2.53	2.66	2.80	2.97	3.09	3.18	3.23	3.34	3.43	3.43
Sichuan	2.27	2.39	2.52	2.65	2.80	2.91	2.99	3.04	3.14	3.23	3.23
Yunnan	2.21	2.13	2.06	1.99	1.91	1.86	1.82	1.80	1.76	1.73	1.73
Guizhou	1.74	1.83	1.91	2.00	2.10	2.17	2.22	2.26	2.32	2.38	2.38

Notes. This table is based on model 4 in table 3 of the main text.

Table A2. Dry land crop yields (dan/mu).

	1653	1692	1729	1766	1808	1836	1856	1868	1892	1911	1912
Shandong	0.43	0.46	0.49	0.52	0.56	0.58	0.60	0.61	0.64	0.66	0.66
Zhili	0.34	0.36	0.38	0.40	0.43	0.45	0.46	0.47	0.48	0.50	0.50
Shanxi	0.25	0.27	0.28	0.29	0.31	0.32	0.33	0.37	0.35	0.36	0.36
Shaanxi	0.52	0.56	0.60	0.64	0.68	0.72	0.74	0.76	0.79	0.82	0.82
Gansu	0.78	0.57	0.43	0.32	0.23	0.19	0.16	0.15	0.12	0.11	0.10
Xinjiang	0.71	0.77	0.83	0.89	0.96	1.02	1.06	1.08	1.14	1.18	1.18
Liaoning	0.29	0.31	0.32	0.34	0.36	0.37	0.38	0.39	0.40	0.41	0.41
Henan	0.73	0.69	0.66	0.62	0.57	0.56	0.55	0.54	0.52	0.51	0.50
Jilin	0.37	0.39	0.41	0.44	0.47	0.49	0.50	0.51	0.53	0.55	0.55
Heilongjiang	0.76	0.82	0.88	0.95	1.03	1.09	1.13	1.16	1.22	1.26	1.27

Notes. This table is based on model 5 in table 4 of the main text.

Bibliography.

Barro, Robert J. (2013), "Inflation and economic growth", *Annals of economics and finance*, vol. 14(1), 85-109.

Bennet, M. K. (1937), "Trends of yield in major wheat regions since 1885", *Wheat studies of the food research institute*, vol. 14(3), 69-102.

Bringas Gutiérrez, Miguel Ángel (2000), *La productividad de los factores en la agricultura Española, 1752-1935* (Madrid: Bank of Spain series *Estudios de historia económica*, no. 39).

Broadberry, Steve, Hanhui Guan and David Daokui Li (2017), "China, Europe and the great divergence: a study in historical national accounting, 980-1850", Oxford University Discussion Paper in Economic and Social History, no. 155.

Chen, Shuo, and James Kai-sing Kung (2016), "Of maize and men: the effect of a NewWorld crop on population and economic growth in China", *Journal of economic growth*, vol. 21(1), 71-99.

Clark, Gregory (2002), "Land rental values and the agrarian economy: England and Wales, 1500-1912", *European Review of Economic History*, vol. 6, pp. 291-308.

Geertz, Clifford (1963), *Agricultural involution: the processes of ecological change in Indonesia* (Berkeley: University of California Press).

Government of Ireland – Central Statistics Office (1997), *Farming since the Famine: Irish farm statistics, 1847-1996*.

Grant, Oliver (2005), *Migration and inequality in Germany 1870-1913* (Oxford: University Press).

Hagen, William W. (2000), *Ordinary Prussians* (Cambridge: University Press).
Li, Bozhong (1998), *Agricultural Development in Jiangnan, 1620-1850* (London: MacMillan).
Main, Jackson T. (1985), *Society and economy in colonial Connecticut* (Princeton: University Press).
Nhemachena, Charity, and Johann Kirsten (2017), "A historical assessment of sources and uses of wheat varietal innovations in South Africa", *South African Journal of Science*, vol. 113(3-4), 1-8.
Overton, Mark (1996), *Agricultural revolution in England* (Cambridge: University Press).
Perkins, Dwight H. (1969), *Agricultural development in China, 1368-1968* (Chicago: Aldine).
Petty, William (1691), *The political anatomy of Ireland* (reprinted Shannon, 1970).
Shi, Zhihong (2017), *Agricultural development in Qing China* (Leiden: Brill).
Simpson, James (2000), "Spanish agriculture in the long run, 1760-1936", Universidad Carlos III, Department of Economic History and Institutions, working paper no. 00-73.
Titow, Jan (1972), *Winchester yields: a study in medieval agricultural productivity* (Cambridge: University Press).
Toutain, Jean-Claude (1961), *Le produit de l'agriculture française de 1700 à 1958*, 2 vols. (Paris: *Cahiers de l'Institut de Sciences Economique Appliquée*, série AF, nos. 1 and 2).
Van Duin, Pieter, and Robert Ross (1987), *The economy of the Cape Colony in the eighteenth century* (Leiden: Centre for the History of European Expansion).

Issued in the series Discussion Papers 2017

2017

01/17 January, **Agnar Sandmo**, "Should the marginal tax rate be negative? Ragnar Frisch on the socially optimal amount of work"

02/17 February, **Luca Picariello**, "Organizational Design with Portable Skills"

03/17 March, **Kurt R. Brekke**, Tor Helge Holmås, Karin Monstad og Odd Rune Straume, "Competition and physician behaviour: Does the competitive environment affect the propensity to issue sickness certificates?"

04/17 March, **Mathias Ekström**, "Seasonal Social Preferences".

05/17 April, Orazio Attanasio, Agnes Kovacs, and **Krisztina Molnar**: "Euler Equations, Subjective Expectations and Income Shocks"

06/17 April, **Alexander W. Cappelen**, Karl Ove Moene, Siv-Elisabeth Skjelbred, and **Bertil Tungodden**, "The merit primacy effect"

07/17 May, **Jan Tore Klovland**, "Navigating through torpedo attacks and enemy raiders: Merchant shipping and freight rates during World War I"

08/17 May, **Alexander W. Cappelen**, Gary Charness, **Mathias Ekström**, Uri Gneezy, and **Bertil Tungodden: "**Exercise Improves Academic Performance"

09/17 June, **Astrid Kunze**, "The gender wage gap in developed countries"

10/17 June, **Kristina M. Bott, Alexander W. Cappelen, Erik Ø. Sørensen** and **Bertil Tungodden**, "You've got mail: A randomized field experiment on tax evasion"

11/17 August, Marco Pagano and **Luca Picariello**, "Talent Discovery, Layoff Risk and Unemployment Insurance"

12/17 August, Ingrid Kristine Folgerø, **Torfinn Harding** and Benjamin S. Westby, «Going fast or going green? Evidence from environmental speed limits in Norway"

13/17 August, **Chang-Koo Chi**, Pauli Murto, and Juuso Välimäki, "All-pay auctions with affiliated values"

14/17 August, **Helge Sandvig Thorsen**, "The effect of school consolidation on student achievement".

15/17 September, Arild Sæther, "Samuel Pufendorf and Ludvig Holberg on Political Economy".

16/17 September**, Chang-Koo Chi,** Pauli Murto, and Juuso Välimäki**, "**War of attrition with affiliated values".

17/17 September, **Aline Bütikofer** and Giovanni Peri, "The Effects of Cognitive and Noncognitive Skills on Migration Decisions"

18/17 October, **Øivind Schøyen**, "What limits the powerful in imposing the morality of their authority?

19/17 October, **Charlotte Ringdal** and **Ingrid Hoem Sjursen**, "Household bargaining and spending on children: Experimental evidence from Tanzania"

20/17 December, **Fred Schroyen** and Karl Ove Aarbu, "Attitudes towards large income risk in welfare states: an international comparison"

21/17 December, **Alexander W. Cappelen, Ranveig Falch**, and **Bertil Tungodden**, "The Boy Crisis: Experimental Evidence on the Acceptance of Males Falling Behind

2018

01/18 January, Øystein Foros, Mai Nguyen-Ones, and **Frode Steen**, "Evidence on consumer behavior and firm performance in gasoline retailing"

02/18 January, **Agnar Sandmo**, "A fundamental externality in the Labour Market? Ragnar Frisch on the socially optimal amount of work"

03/18 February, Pierre Dubois and **Morten Sæthre, "**On the Effect of Parallel Trade on Manufacturers' and Retailers' Profits in the Pharmaceutical Sector"

04/18 March, **Aline Bütikofer**, Julie Riise, and Meghan Skira, "The Impact of Paid Maternity Leave on Maternal Health"

05/18 March, **Kjetil Bjorvatn** and **Bertil Tungodden**, "Empowering the disabled through savings groups: Experimental evidence from Uganda"

06/18 April, Mai Nguyen-Ones and **Frode Steen,** "Measuring Market Power in Gasoline Retailing: A Market- or Station Phenomenon?

07/18 April, **Chang Koo Chi** and Trond Olsen, "Relational incentive contracts and performance"

08/18 April, **Björn Bartling**, **Alexander W. Cappelen, Mathias Ekström, Erik Ø. Sørensen,** and **Bertil Tungodden**, «Fairness in Winner-Take-All Markets»

09/18 April, **Aline Bütikofer, Sissel Jensen**, and **Kjell G. Salvanes**, «The Role of Parenthood on the Gender Gap among Top Earners»

10/18 May, **Mathias Ekström**, "The (un)compromise effect"

11/18 May, Yilong Xu**, Xiaogeng Xu**, and Steven Tucker**,** «Ambiguity Attitudes in the Loss Domain: Decisions for Self versus Others»

12/18 June, **Øivind A. Nilsen**, Per Marius Pettersen, and Joakim Bratlie, "Time-Dependency in producers' price adjustments: Evidence from micro panel data"

13/18 June, **Øivind A. Nilsen**, Arvid Raknerud, and Diana-Cristina Iancu, "Public R&D support and firms' performance. A Panel Data Study"

14/18 June, Manudeep Bhuller, Gordon B. Dahl, **Katrine V. Løken**, and Magne Mogstad: «Incarceration, Recidivism, and Employment"

15/18 August, Manudeep Bhuller, Gordon B. Dahl, **Katrine V. Løken**, and Magne Mogstad: «Incarceration Spillovers in Criminal and Family Networks"

16/18 August, Pedro Carneiro, Kai Liu, and **Kjell G. Salvanes**: "The Supply of Skill and Endogenous Technical Change: Evidence From a College Expansion Reform"

17/18 August, **Chang Koo Chi**, "An analysis of the two-bidder all-pay auction with common values"

18/18 August**, Alexander W. Cappelen,** Cornelius Cappelen**,** and **Bertil Tungodden, "**Second-best fairness under limited information: The trade-off between false positives and false negatives"

19/18 September, **Aline Bütikofer, Antonio Dalla Zuanna**, and **Kjell G. Salvanes**: "Breaking the Links: Natural Resource Booms and Intergenerational Mobility"

20/18 September, Juan Pablo Atal, José Ignacio Cuesta, and **Morten Sæthre**, "Quality regulation and competition: Evidence from Pharmaceutical Markets"

21/18 October, Orazio Attanasio, Agnes Kovacs, and **Krisztina M**olnar, "Euler Equations, Subjective Expectations and Income Shocks"

22/18 October, Antonio Mele, **Krisztina Molnár**, and Sergio Santoro, "On the perils of stabilizing prices when agents are learning"

23/18 November, Bjørn-Atle Reme, Helene Lie Røhr, and **Morten Sæthre**, "The Poking Effect: Price Changes, Information, and Inertia in the Market for Mobile Subscriptions"

24/18 November, **Ingrid Hoem Sjursen**, "Accountability and taxation: Experimental evidence"

25/18 November, **Liam Brunt** and Antonio Fidalgo, "Why 1990 international Geary-Khamis dollars cannot be a foundation for reliable long run comparisons of GDP"

26/18 November, **Ola Honningdal Grytten**, "A continuous consumer price index for Norway 1492-2017"

27/18 December, **Liam Brunt** and Antonio Fidalgo "Feeding the people: grain yields and agricultural expansion in Qing China"

Norges
Handelshøyskole

Norwegian School of Economics

NHH
Helleveien 30
NO-5045 Bergen
Norway

Tlf/Tel: +47 55 95 90 00
Faks/Fax: +47 55 95 91 00
nhh.postmottak@nhh.no
www.nhh.no