Essays in Labor and Health Economics

By

Timothy M. Watts

B.A. Economics
Vanderbilt University, 2002

SUBMITTED TO THE DEPARTMENT OF ECONOMICS IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN ECONOMICS
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SEPTEMBER 2006
[September 2007]

Signature of Author: ____________________
Department of Economics
August 15, 2007

Certified by: ____________________
David Autor
Professor of Economics

Certified by: ____________________
Joshua Angrist
Professor of Economics

Accepted by: ____________________
Peter Temin
Professor of Economics
Chairman, Committee for Graduate Students

Essays in Labor and Health Economics

By

Timothy M. Watts

Submitted to the Department of Economics on August 15, 2007 in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Economics

Abstract

This dissertation consists of three essays in empirical labor and health economics.

The first chapter examines how the amount of time devoted to a leisure activity varies in response to temporary changes in the price of that activity. Specifically, I estimate the effect of changes in expected winnings in an online poker game on the probability that players quit playing. I find that expected winnings have a large negative effect on the probability that a player quits playing poker. A one dollar increase in expected winnings decreases the probability that a player quits playing altogether by 0.5 percentage points, compared to the mean of 1.1 percentage points. This corresponds to a price elasticity of demand for poker of -0.14.

The second chapter develops and tests a model of how college students choose their field of study. The model combines features from learning and human capital models and captures several stylized facts from the empirical literature on choice of college major. I test the model's predictions using High School and Beyond data. I find three results that generally agree with the model's predictions. First, students with higher levels of ability choose majors with higher average earnings. Second, students who receive low grades in college are more likely to change their field of study. Third, students who switch majors in college subsequently earn less than students who do not change majors, but this difference is primarily due to major-switchers obtaining degrees in low-paying fields.

The third chapter, coauthored with Abhijit Banerjee, Esther Duflo and Gilles Postel-Vinay, provides estimates of the long-term effects on height and health of a large income shock experienced in early childhood. Phylloxera, an insect that attacks the roots of grape vines, destroyed 40% of French vineyards between 1863 and 1890, causing major income losses among wine growing families. We exploit the regional variation in the timing of this shock to identify its effects. We find that, at age 20, those born in affected regions were about 1.8 millimeters shorter than others. This estimate implies that children of wine-growing families born when the vines were affected in their regions were 0.6 to 0.9 centimeters shorter than others by age 20. This is a significant effect since average heights grew by only 2 centimeters in the entire 19th century.

Thesis Supervisor: David Autor
Title: Associate Professor of Economics

Thesis Supervisor: Joshua Angrist
Title: Professor of Economics

Acknowledgements

My thesis advisors, Joshua Angrist and David Autor, supplied much needed inspiration, direction and constructive criticism for this research. I am greatly indebted to both of them. Esther Duflo also provided valuable feedback on my research, and I learned many practical lessons from working with her as a coauthor.

I thank the National Science Foundation Graduate Research Fellowship for providing financial support during my graduate studies.

My classmates provided much help in the completion of this project. I would especially like to thank Peter Hinrichs, Jin Li and Peter Schnabl, who generously shared their time and knowledge. I also thank Michelle Liu for her counsel and companionship.

Finally I thank my family for their unfailing encouragement and particularly my parents, Michael and Mary, who have continually demonstrated their willingness to aid my endeavors, academic and otherwise.

Table of Contents

The Wages of Sin: Leisure Choices of Online Poker Players

Timothy M. Watts

Abstract

This paper examines how the amount of time devoted to a leisure activity varies in response to temporary changes in the price of that activity. Specifically, I estimate the effect of changes in expected winnings in an online poker game on the probability that players quit playing. I estimate this effect with data on over 300,000 hands of online poker, using a measure of opponent skill as an instrument for expected winnings. I find that expected winnings have a large negative effect on the probability that a player quits playing at his current table. A one dollar increase in expected winnings decreases the probability of leaving the current table by 0.9 percentage points, compared to the mean probability of 2.4 percentage points. After players leave one table, they can start playing at another one immediately, so substitution between poker tables rather than between poker and other activities accounts for some of the effect of expected winnings. Even so, expected winnings affects the total amount of poker played. A one dollar increase in expected winnings decreases the probability that a player quits playing altogether by 0.5 percentage points, compared to the mean of 1.1 percentage points. This corresponds to a price elasticity of demand for poker of -0.14. The results indicate that poker players respond to changes in the price of poker by playing more when the price of poker is low.

I. Introduction

People devote a large portion of their time to leisure activities. Aguiar and Hurst (2006) report that Americans spend 36 hours a week in leisure activities compared to 51 hours in market and nonmarket work. In theory, changes in the prices of leisure activities should affect the consumption of leisure activities in the same way that changes in wages affect labor supply. Yet while an extensive literature estimates the relationship between wages and labor supply, few papers estimate the relationship between the price and consumption of leisure.

This chapter examines how temporary changes in the price of a leisure activity affect the amount of time devoted to that activity. Specifically, I estimate the effect of changes in expected poker winnings on the probability that players quit playing. Because poker players lose money on average, expected winnings will typically be negative and can be considered the price of poker. Expected winnings affect the amount of poker played through two channels. First, an increase in expected winnings makes poker less expensive compared to other activities, which causes players to allocate more time to poker. Second, an increase in expected winnings increases expected wealth, but whether this increases or decreases the amount of poker played depends on the player's utility function.

To develop these ideas, I model the poker-playing decision as the solution to a dynamic income- and time-allocation problem. Within a period, players allocate income between consumption and poker and allocate time between poker and some costless leisure activity. The model shows that, as long as wealth effects are small, an increase in expected winnings decreases the likelihood that a player quits playing poker.

I test this prediction using data on over 300,000 hands of online poker. From this data I construct records of poker sessions, which are groups of hands played by a single player at the same table and seat without long gaps between consecutive hands. The sample yields over 55,000 sessions played by more than 4,500 players. With this data I estimate the relationship between expected winnings and quitting. Expected winnings are unobserved, so I use actual winnings as a proxy. A regression of an indicator variable for the last hand of a session on actual winnings suggests that actual winnings and, by implication, expected winnings have a very small negative effect on the probability of ending a session. These estimates are biased towards zero, however, since actual winnings are a noisy measure of expected winnings.

To correct the attenuation bias, I use a measure of opponent skill as an instrumental variable for expected winnings. The measure of opponent skill is observable by players, and anecdotal evidence suggests players pay attention to this information. While players can choose the opponents they face at the beginning of the session, within-session opponent turnover creates exogenous variation in opponent skill. I use this within-session variation in opponent skill to identify the effect of expected winnings.

The instrumental variable estimates suggest that expected winnings have a strong negative effect on the probability of ending a session, which occurs when a player stops playing poker at his current table. A one dollar increase in expected winnings decreases the probability of ending a session by 0.9 percentage points, which is large compared to the mean probability of 2.4 percent. Part of the effect of expected winnings on session length is due to players switching between poker tables rather than players switching between poker and some other activity. Yet further investigation reveals that expected winnings affect the total amount of poker played as well. A one dollar increase in expected winnings reduces the probability that a player quits playing poker altogether by 0.5 percentage points. From this estimate, I calculate that the price elasticity of poker consumption is -0.14. Together these results show that poker players respond to changes in the price of poker by playing more when the price of poker is low.

Two recent papers find that people take less leisure when the weather is bad. Neidell (2006) finds that participation in certain outdoor leisure activities in Southern California decreases on days with smog alerts. Connolly (2006) finds that men spend more time at work and less time on leisure on rainy days. Since bad weather likely reduces the marginal utility of leisure, these results suggest a positive relationship between the marginal utility of leisure and the amount of leisure consumed. But without further research on the willingness to pay to avoid bad weather, these results cannot be interpreted in monetary terms. This paper builds on the literature by examining the response of leisure time to the price of a leisure activity. My estimates of the effect of expected winnings on the probability of ending a poker session translate directly into a price elasticity of poker. Thus the leisure demand response I estimate may be directly compared with estimates of labor supply elasticity.

This paper is also methodologically and theoretically related to several recent studies on labor supply, which use high frequency data from markets in which the workers have flexible hours and variable wages. Several of these papers examine the labor supply of taxi drivers. Camerer et al. (1997) and Chou (2000) find that taxi drivers work fewer hours on days with higher average wages. The authors suggest that this evidence supports the target-earner model of labor supply, in which workers stop working once they have reached some target level of income. Farber (2005), however, estimates a survival time model of taxi driver labor supply and finds that daily income does not predict the probability that a driver stops working. Other recent studies find that bicycle messengers work more days when wages are high (Fehr and Götte, 2005) and that stadium vendors are more likely to work at games that have high expected wages (Oettinger, 1999). My paper contributes to this literature by using a similar strategy to estimate the price elasticity of demand for leisure.

The paper is also related to a large empirical literature on gambling. One main strand of the gambling literature tests the efficiency of various wagering markets (see Sauer, 1998, for a survey). Although studies find some persistent deviations from efficient pricing, the main result is that wagering markets are approximately efficient. A second focus of the literature is the demand for lottery tickets (see e.g. Guryan and Kearney, 2005; Kearney, 2006; and Clotfelter and Cook, 1993). In contrast to much of the empirical gambling literature, which uses data on aggregate variables, this paper uses individual-level data. My paper contributes to the empirical gambling literature by offering evidence that one aspect of gambling behavior is consistent with utility maximization.

II. Background on Poker

Poker has surged in popularity in the past three years. The World Series of Poker, an annual poker tournament with a $10,000 entry fee, provides a striking illustration. Between 2003 and 2006, the number of World Series entrants increased tenfold – from 839 to 8,773. The online poker industry has also grown rapidly. The first online poker site opened in 1998. Today hundreds of poker sites cater to millions of players and the earnings and volume of online poker exceed that of brick and mortar establishments[1].

Online card rooms offer several varieties of poker. This paper focuses on one particular poker variant called fixed-limit Texas hold'em. Like all poker games, fixed-limit Texas hold'em is a strategic game of imperfect information. Players vie against one another to win the money in the pot, which contains the wagers made by the players in the course of a game. A single game of fixed-limit Texas hold'em, called a hand, consists of two to ten players and up to four rounds of betting. Before the first round of betting each player receives two cards facedown and two players must make forced bets. Play proceeds clockwise around the table beginning with the player to the left of the forced bets. To continue playing, players must either match or exceed whatever bets have been made in the current round. In fixed-limit Texas hold'em poker, players may only bet a predetermined amount of money. If a player does not match a previous bet, he gives up his cards and forfeits the hand.

At the end of the first, second, and third betting rounds, additional cards are dealt face up in the middle of the table. A hand ends when a player makes a bet that no other player matches or at the end of the fourth betting round. If a player makes a bet that no one matches, or if he makes the best poker hand at the end of the final betting round, he wins all the bets.

III. A Model of Demand for Poker

Online poker rooms operate continuously, and they increase the number of virtual poker tables to accommodate new players. Thus a player may join an online poker game at any

[1] Online poker sites earned \$3 billion in 2005 (Xanthopoulos, 2006), compared to \$1.2 billion for card rooms and casino poker (AGA, 2006). This suggests that most poker hands take place online, since online poker sites typically earn less money per hand than card rooms.

time and, once he has joined, he may play as much or as little poker as he likes. The end of each hand provides the player with an opportunity to decide whether to continue playing by comparing the costs and benefits of one more hand of poker. Thus a complete model of the demand for poker would relate a player's decision after each hand to variables such as how long he has been playing and how much he expects to win or lose in the next hand. Because expected winnings may predictably rise and fall over the course of many hands, the solution to the full model is difficult to characterize. I follow the approach of Farber (2005) and develop a simpler model whose solution gives a reasonable approximation to the more complicated optimal stopping problem.

A. Assumptions

Each period the player derives utility from the consumption of goods, C_t, and time spent playing poker, H_t. On average, the player loses money at poker, so he allocates wealth between goods and poker. The player divides his time between poker and some other costless leisure activity[2]. The player's utility function is additively separable between periods and is defined for period t as

(1) $$U_t = \sum_{s=t}^{T} E[U(C_s, H_s)].$$

The period utility function is twice differentiable and strictly concave. Another important assumption is that poker winnings do not affect the marginal period utility of poker.

The player has initial assets A_0 and accumulates assets according to:

(2) $$A_{t+1} = A_t + w_t H_t - C_t.$$

[2] Adding a work activity to the model results in the same predictions, but I omit work since I have no data on the (non-poker) labor supply and wages of the players in my sample.

w_t is winnings, which is the average amount of money won per unit of time spent playing poker in period t. In words, the player's assets in the next period depend on assets at the beginning of the current period, the amount of won or lost playing poker and expenditure on consumption goods. During a period, the player plays poker against a fixed set of opponents. At the beginning of period t, the player observes the skill of his opponents without error and correctly forecasts w_t from this information[3]. The set of opponents varies between periods, and future values of w are uncertain. Since w_t is on average negative, $-w_t$ can be thought of as the price of poker, because it represents the amount of consumption a player must forgo to play one unit of poker.

B. The Player's Maximization Problem

Each period the player chooses C_t and H_t to maximize the value function

(3) $$V(A_t,t) = (U(C_t,H_t) + E[V(A_{t+1},t+1)]),$$

subject to the asset accumulation rule (2). The optimal choice of consumption satisfies

(4) $$U_C(C_t,H_t) = \lambda_t \equiv \partial V / \partial A_t,$$

which means the marginal utility of consumption equals the marginal utility of wealth. The optimal choice of poker satisfies

(5) $$U_H(C_t,H_t)/\lambda_t \geq (-w_t),$$

which means the ratio between the marginal utililty of poker and marginal utility of wealth is greater than or equal to ratio between the price of poker and the price of consumption (which has been normalized to one). Finally, the allocation of wealth across periods is determined by

[3] In the model, the player knows the winnings per poker unit with certainty. This differs from reality since players do not know the winnings from any hand of poker in advance. The assumption may be justified, however, if players play enough hands in each period that the variance in actual per hand winnings is arbitrarily small.

(6) $\lambda_t = E_t[\lambda_{t+1}]$,

which means the marginal utility of wealth in the current period equals the expected marginal utility of wealth in the next period. The terminal condition

(7) $A_{T+1} = 0$

completes the characterization of the player's solution.

C. Theoretical Predictions

Equation (5) implicitly defines H_t as a function of C_t, w_t and λ_t. Let h_t denote a single hand of poker in period t, and let $\bar{h}_t$ denote the final hand of poker in period t. The probability that a player quits playing poker after a given hand can be written as

(8) $p_{tm} = \Pr[h_t = \bar{h}_t \mid w_t, H_t] = \Phi(-U_H(C_t, H_t)/\lambda_t - w_t)$,

where H_t is the amount of poker played so far in the period and $\Phi(.)$ is a monotonic increasing function. Equation (8) makes two predictions.

First, each hand of poker played increases the probability of stopping. This follows from the assumption that $U_{HH} < 0$.

Second, an increase in winnings typically decreases the probability of stopping. Winnings influence p through a substitution effect and a wealth effect. The substitution effect is unambiguously negative. The sign of the wealth effect depends on the sign of U_H. As long as U_H is positive, then the wealth effect of an increase in winnings is negative and reinforces the substitution effect. But if U_H is negative, the wealth effect offsets the substitution effect. In either case, as long as winnings are small compared to total wealth, the wealth effect should also be small.

IV. Empirical Strategy

A. Estimating Equations

The theory suggests that an increase in expected winnings decreases the probability that a player stops playing poker. I test this prediction empirically by comparing the quitting behavior of poker players across hands with different expected winnings. The unit of observation is a player-hand, and each player-hand belongs to a player-session. If I define $\Phi(.)$ from equation (8) as $\Phi(x) = x$ the model

$$(9) \qquad M_{hti} = \alpha_{ti} + \beta_1 w_{hti} + \beta_2 h_{hti} + \beta_3 X_{hti} + \varepsilon_{hti}$$

gives a linear approximation of the relationship between the probability of quitting and control variables. The dependent variable M_{hti} is an indicator variable that equals one if hand h is the final hand of session t for player i. α_{ti} is a session fixed effect, w_{hti} is the expected winnings in dollars for the next hand, X_{hti} is a vector of effects for the number of players in the hand and ε_{hti} is the error term. The coefficient β_1 measures the change in probability of quitting caused when expected winnings for the next hand increase by one dollar. With session effects included, within-session variation in expected winnings identifies β_1.

The econometric model can be viewed as a reduced-form linearization of equation (8). The probability of quitting should be negatively related to expected winnings and positively related to the number of hands played. Thus the theoretical prediction is that $\beta_1 < 0$ and $\beta_2 > 0$.

Since expected winnings are unobservable, β_1 cannot be estimated using equation (9). I estimate instead

$$(10) \qquad M_{hti} = \alpha_{ti} + \beta_1 \text{ActualWinnings}_{hti} + \beta_2 h_{hti} + \beta_3 X_{hti} + \varepsilon_{hti},$$

where $\text{ActualWinnings}_{hti}$ proxies for w_{hti}. $\text{ActualWinnings}_{hti}$ is defined as the net dollar amount won by player i in hand h of session t. $\text{ActualWinnings}_{hti}$ equals expected winnings plus measurement error. Since the amount won in a given hand depends largely on chance, the error term is large. Thus estimating (10) by ordinary least squares (OLS) produces biased estimates. Specifically, $\hat{\beta}_1$ understates the magnitude of the effect of expected winnings on the probability of quitting.

To correct for this attenuation bias, I estimate equation (10) by two stage least squares (2SLS), using a measure of opponent skill as an instrumental variable for expected winnings. The first stage relationship between actual winnings and opponent skill is

(11) $$\text{ActualWinnings}_{hti} = \delta_{ti} + \gamma_1 \text{OpponentSkill}_{hti} + \gamma_2 h_{hti} + \gamma_3 X_{hti} + u_{hti},$$

where δ_{ti} is a session fixed effect and u_{hti} is the error term.

B. Threats to Validity

Whether estimating equation (10) by 2SLS produces consistent coefficients depends on three conditions. First, my measure of opponent skill must be correlated with actual winnings. Estimates of the first stage equation (11) in section VI show that this condition is satisfied.

Second, opponent skill must be uncorrelated with the error term of equation (10). If opponent skill instead affects the probability of quitting through an omitted variable, the 2SLS estimates will be inconsistent. Two possible concerns are that opponent skill is correlated with player-specific characteristics or the marginal utility of poker, both of which affect the probability of quitting. For example, players who like to play long sessions of poker may also select tables with worse opponents. But including session effects addresses this concern because they control for fixed player characteristics. A greater concern is that players may enjoy playing poker against certain types of opponents. If players enjoy playing against unskilled opponents, the 2SLS estimates will overstate the relationship between expected winnings and the probability of quitting. On the other hand, players may enjoy the challenge of playing against skilled opponents, which would bias the estimates in the opposite direction. Section VII investigates these concerns further.

Third, opponent skill must be uncorrelated with the difference between expected and actual winnings. Otherwise, the first-stage projections of opponent skill onto actual winnings capture expected winnings plus measurement error, and the 2SLS estimates will be biased toward zero. In the model, opponent skill is uncorrelated with the measurement error by assumption since players form expected winnings conditional on opponent skill.

C. *Measuring Opponent Skill*

I use the poker statistic VP$IP to measure opponent skill. Voluntarily puts money in the pot (VP$IP) is defined as the fraction of hands for which a player calls or places a bet in the first round of betting, excluding any forced bets. The theoretical relationship between VP$IP and win rates is clear. To maximize winnings, players should only put money in the pot when they expect to win as much or more money in return[4]. The amount of money a player expects to win depends in part on the probability that his cards will make the best poker hand. If players consistently call bets with cards that have little chance of making the best hand or consistently fold cards that have a strong chance of winning, they will lose money. This suggests a nonmonotonic relationship between VP$IP and skill, although in practice players are far more likely to play too many hands than too few. Thus VP$IP is likely inversely related to skill.

[4] This ignores any effects of the current hands on the play of future hands.

To construct my instrument, for each player-hand I calculate the VP$IP of each opponent using all hands in my sample except for the hands that are part of the same player-session. For example, if player A plays a session that last for 40 hands, the VP$IP for all of player A's opponents will be calculated excluding the 40 hands from that session[5]. This method prevents mechanical correlation between the opponents' VP$IPs and a player's actual winnings. The opponents' VP$IPs are then averaged. The resulting measure, called opponent VP$IP, remains constant throughout a player-session unless the composition of opponents changes.

V. Data

To estimate equation (10), I collected data from detailed poker hand records from an online card room. All hands in the sample are fixed-limit Texas hold'em at $60 stakes and took place between January 9th and 26th, 2006. The sample contains 320,009 hands with an average of 7.4 players per hand. Each hand's record contains its time, date and table number as well as the usernames, betting actions and net winnings of each player involved. Each player is identified by a unique username, which allows me to track a player's behavior over many hands[6].

A. Player Statistics

Panel A of table 1 shows descriptive statistics for players. The sample contains 4,586 players. Players play, on average, 12 sessions, 519 hands and 391 minutes of poker in the sample. There is substantial variation in these measures, and the mean is greater than the

---

[5] Some player-sessions contain every hand of an opponent with few hands in the sample. In this situations, no VP$IP is calculated for that opponent and mean VP$IP is the average of the remaining opponents' VP$IPs. A small number of player-hands are excluded from the analysis because no opponent's VP$IP can be calculated.

[6] Players may change their usernames once every six months, so a small portion of players in my sample could have two usernames. If a player abandons a username, no other player may use it.

median for each one, which suggests a right-skewed distribution of poker consumption in which a few players consume large amounts of poker.

Players lose an average of $148 in the sample at the rate of $571 per hour or $7.20 per hand. Again, there is substantial variation in losses. 39 percent of players have positive winnings. The 25th percentile of winnings per hand is $-12, the median is $-2, and the 75th percentile is $3. The medians of total winnings and win rates exceed their means, suggesting a left-skewed distribution of winnings where a few players lose much more than the average amount.

The average player's VP$IP is 0.37, which means the average player risks money voluntarily with 37 percent of hands. The standard deviation of player VP$IP is 0.14, which suggests poker strategy differs widely between players.

B. Session Statistics

Each hand belongs to a session – defined as a series of hands played by a single player at the same seat and table with a maximum gap of 30 minutes between consecutive hands. Panel B shows descriptive statistics of sessions. The sample contains 55,983 player-sessions. Sessions last on average for 43 hands or 32 minutes, and the standard deviation of session length is 51 hands or 39 minutes. Between-player differences likely explain some of the variation in session length. Among players with at least two sessions, however, the average *within*-player standard deviation of session length is 34 hands or 26 minutes. This within-player dispersion in session length could result from short-term variation in expected winnings. Players lose $12.15 in an average session.

C. Player-hand Statistics

The sample contains 2.38 million player-hands, which is the unit of observation for the econometric model. Panel C shows summary statistics for the variables used in the regressions. The sample average of the primary dependent variable, a dummy for the last hand of a session, is 0.024, which means the unconditional hazard of ending a session is 2.4 percent. Players lose $0.30 in an average player-hand[7]. The mean value of opponent VP$IP, the instrumental variable for expected winnings, is 0.28 with a standard deviation of 0.05.

VI. Empirical Results

A. Relationship between Expected Winnings and the Probability of Quitting

I estimate the relationship between the probability of ending a poker session and expected winnings, using opponent skill as an instrument for expected winnings. The data suggest a predominantly negative relationship between a player's VP$IP and his own winnings. Figure 1 shows the results of a kernel regression of a player's winnings per hand on his VP$IP. Winnings initially increase in VP$IP and then trend steadily downward. The maximum predicted winnings occur when VP$IP equals 0.16, but most players in the sample exceed this VP$IP: the median VP$IP is 0.34, the 25th percentile is 0.24 and the 7th percentile is 0.16. Thus the typical player's expected winnings should decrease in his own VP$IP and increase in the VP$IP of his opponents.

Column 1 of table 2 shows the first-stage relationship between opponent VP$IP and actual winnings. As expected, players win more against opponents with higher VP$IPs. A one percentage point increase in opponent VP$IP causes a $0.21 increase in actual winnings. The t-statistic of this estimate is 8.2. Because there may be serial

[7] This number is negative because the card room generates revenue by collecting a portion of the money wagered in each hand. Per hand winnings are greater when averaged across player-hands than when averaged across players because a player's total hands are positively correlated to his per hand winnings.

correlation in opponent VP$IP, standard errors allow for clustering within a player-session. The regression includes session fixed effects, the number of hands played so far in the session, and a set of dummies for the number of players that participated in the hand.

Column 2 of table 2 shows the OLS relationship between the probability of quitting and actual winnings, which proxy for expected winnings. The coefficient on net dollars won is -0.0021, which corresponds to the percentage point change in the probability of quitting for a $1 increase in net winnings. Taken at face value, this result suggests that expected winnings have a small negative effect on the probability of quitting. Yet the actual amount won in a hand contains expected winnings plus noise, so the OLS estimate will be biased towards zero. The inclusion of session effects exacerbates this bias since expected winnings are likely serially correlated within a session. One way to reduce the attenuation bias is to use a proxy with less noise, such as the average of actual winnings over several hands. Column 3 shows the OLS estimate using cumulative session winnings per hand as a proxy. The coefficient on actual winnings is -0.0024, which is still small and not significantly different from the estimate of column 2.

Column 4 shows the reduced-form relationship between the probability of quitting and opponent VP$IP. The results suggest that players play longer against less skilled opponents: a one percentage point increase in opponent VP$IP decreases the probability of quitting by 0.19 percentage points. Column 5 shows the 2SLS estimate of the relationship between the probability of quitting and expected winnings. Because opponent VP$IP should be uncorrelated with the difference between actual and expected winnings, using it as an instrument for expected winnings should correct for the attenuation bias of columns 2 and 3. Players quit playing less often when expected winnings are high. Increasing expected winnings by one dollar decreases the probability of quitting by 0.9 percentage points. This is a large effect compared to the mean probability of quitting of 2.4 percent. The 2SLS estimate is several orders of magnitude larger than the OLS estimate, which suggests the attenuation bias of OLS is severe. This is no surprise since the actual amount won in a given hand depends on many random variables besides expected winnings, such as a player's cards and the cards of each of his opponents.

As predicted by the model, the probability of quitting is negatively related to expected winnings and positively related to the number of hands played. This implies that the length of a poker session is also negatively related to expected winnings. To quantify this relationship, I calculate the elasticity of session length with respect to expected winnings using the coefficients in table 2. Since the coefficient on hands played is small, and the typical session consists of only 43 hands, a constant hazard function should allow a reasonable approximation of expected session length. Define the constant hazard function as

$$\lambda(h) = C + \hat{\beta}_1 w, \tag{12}$$

where the constant C is defined as the difference between the sample mean probability of quitting and the product of $\hat{\beta}_1$ and the sample mean of actual winnings. $\lambda(h)$ gives the probability of quitting for a given value of expected winnings, holding other variables fixed at their sample means.

For a constant hazard, the reciprocal of the hazard gives the expected duration until failure. The elasticity of the expected number of hands remaining in the session with respect to expected winnings is

(13) $\quad \eta = -\hat{\beta}_1 w/(C + \hat{\beta}_1 w)$.

Using the sample mean value of per hand winnings of \$-0.30 for w, the OLS estimates from column 2 correspond to an elasticity of -0.003 while the 2SLS estimates from column 4 correspond to an elasticity of -0.113. Because mean winnings are negative, these should be interpreted as price elasticities of poker consumption[8].

B. Switching Tables versus Logging Off

The results of tables 2 support the hypothesis that a decrease in expected winnings increases the likelihood that a player quits his current session. After a player quits one session, however, he may immediately begin another at a different table. In fact, the results would be consistent with players who play the same amount of poker every day, but switch to a new table when expected winnings fall at the current table. To test whether expected winnings affect the total amount of poker played rather than just inducing players to switch tables, I regress an indicator variable for logging off on winnings and other controls. The indicator equals one if the hand is the last of a session *and* the player's next in-sample session does not begin for at least 15 minutes.

Table 3 shows the relationship between expected winnings and the probability of logging off. Column 1 shows the OLS estimates. A one dollar increase in actual winnings decreases the probability of logging off by 0.0016 percentage points. Column 2 shows the 2SLS estimates. A one dollar increase in expected winnings decreases the

[8] For ease of presentation, the coefficients in the relevant columns in table 2 are reported as percentage point effects. They must be divided by 100 to perform the calculations of equations (12) and (13).

probability of logging off by 0.47 percentage points. These estimates suggest that substitution between tables accounts for half of the effect of expected winnings on the probability of quitting a poker session.

Since at any time my sample contains a maximum of ten tables while the online card room operates up to 25 tables of the same poker type and stakes, I overestimate the number of players who log off. Multiplying the coefficients in columns 1 and 2 by 0.4 gives a conservative correction for the overestimation of the dependent variable. The resulting OLS and 2SLS estimates are 0.0006 and 0.20. This correction is not perfect, however, since the number of active tables varies during the sample. Unfortunately, the data do not give the number of tables active at a given time.

One way to work around this problem is to use only data from times when my sample contains all active tables. Because my data-collection software automatically collected records from up to ten tables, I know that my sample contains all active tables for periods that contain data from nine or fewer tables. Columns 3 and 4 of table 3 restrict the sample to periods in which there were nine or fewer active tables for at least 30 minutes. A one dollar increase in actual winnings decreases the probability of logging off by 0.0018 percentage points and a one dollar increase in expected winnings decreases the probability of logging off by 0.54 percentage points.

Hands in the restricted sample take place during off-peak hours, such as early mornings on weekdays. If players playing at those times have relatively inelastic demand for poker, the results of column 4 would not provide a good estimate of typical behavior. To check the amount of difference between the full and restricted sample, I estimate the relationship between expected winnings and the probability of ending a session using the restricted sample. Columns 5 and 6 report the results. Increasing expected winnings by one dollar decreases the probability of ending a session by one percentage point. The similarity between this estimate and the full-sample estimate suggest that players in the restricted sample behave like the typical player.

The estimates from the restricted sample suggest that substitution between poker tables accounts for half of the effect of expected winnings on the probability of ending a session. They also show, though, that changes in expected winnings affect the total amount of time devoted to poker. Repeating the elasticity calculation from the previous subsection using the estimates from column 4 yields an elasticity of total hands of poker played in a period with respect to expected winnings of -0.136. This is slightly larger than the elasticity calculated in the previous section, which seems puzzling since the effect of expected winnings on the probability of logging off is smaller than the effect on ending a session. But the mean probability of logging off (1.2 percent) is lower than the mean probability of ending a session (2.4 percent) and the effect of expected winnings varies in proportion to the means. The similarity of session-length and period-length elasticities suggests that the total amount of poker played in a given day is proportional to session length.

C. *Estimating Effects by Experience Level*

If players gradually learn how to solve the optimization problem of section III, the effect of expected winnings on the probability of quitting may increase in player experience. Other studies have found that people's behavior differs by their level of experience. Camerer et al. (1997) find that, in two of three samples, high-experience drivers have significantly higher wage elasticities than low-experience drivers. List finds that more experienced market participants bargain better (2004) and are less likely to display the endowment effect (2003).

I divide players into experience groups based on a simple criterion: a player's experience is high if I ever observe him playing at two or more tables simultaneously[9]. Playing at multiple tables requires more effort and concentration and faster decision making than playing at one table. Relatively inexperienced players may find playing at a single table taxing, especially since online play proceeds much faster than live play. Even seasoned players who could handle multiple tables may prefer to play one table since this frees them to perform other tasks at the same time. On the other hand, experienced players and winning players will be more attracted to multi-table play, since they can more easily manage additional tables and playing extra tables will increase their expected hourly winnings.

Table 4 shows summary statistics for high- and low-experience players. Slightly more than one quarter of players are classified as high experience, but because they play nearly 9 times as many hands, sessions and minutes as low-experience players, they account for three quarters of player-hand observations. High-experience players are also more successful. They lose on average only \$0.19 per hand compared with a \$9.79 per hand loss of low-experience players. High-experience players also use different strategies – their VP\$IP is 0.28 compared to 0.40 for low-experience players. In short, high-experience players play much more poker and lose far less than low-experience players.

[9] The data do not contain such natural measures of experience such as the length of time since the player joined the card room or the total number of hands the player has played on the cardroom. My criterion for experience captures a mix of characteristics including skill and effort.

Table 5 shows the results when the regressions from tables 2 and 3 are run separately for players of high and low experience. Column 1 shows the first-stage relationship between opponent VP$IP and dollars won in the current hand. A one percentage point increase in opponent VP$IP increases winnings by $0.43 for inexperienced players and $0.13 for experienced players. Columns 2 and 3 report the estimates for the probability of quitting. Column 2 shows the OLS estimates. A one dollar increase in actual winnings decreases the probability of quitting by 0.004 percentage points for inexperienced players and by 0.001 percentage points for experienced players. Column 3 shows the 2SLS estimates. A one dollar increase in expected winnings decreases the probability of quitting by 0.43 percentage points for inexperienced players and by 1.42 percentage points for experienced players.

Columns 4 and 5 report the estimates for the probability of logging off. Column 4 shows the OLS estimates. A one dollar increase in actual winnings decreases the probability of logging off by 0.0035 percentage points for inexperienced players and by 0.0007 percentage points for experienced players. Column 5 shows the 2SLS estimates. A one dollar increase in expected winnings decreases the probability of logging off by 0.36 percentage points for inexperienced players and by 0.59 percentage points for experienced players.

Table 5 reveals three interesting differences between inexperienced and experienced players. First, the first-stage effect of opponent VP$IP on actual winnings is three times larger for inexperienced player compared to experienced players. One plausible explanation is that inexperienced players use the same strategy against all opponents while experienced players use different strategies against different opponents.

Second, expected winnings have a bigger effect on both the probability of ending a session and the probability of logging off for experienced players. This supports the idea that experienced players' demand for poker responds more strongly to changes in expected winnings. Either a higher elasticity of demand or the ability to more accurately estimate expected winnings or a combination of both could explain why the demand response is greater for experienced players.

Third, substitution between tables (as a proportion of the overall response to changes in expected winnings) is greater for experienced players. Substitution between tables accounts for 60 percent of experienced players' response to expected winnings but only 20 percent of inexperienced players' response. By switching tables relatively more often compared to inexperienced players, experienced players may be able to play against lower skilled opponents on average, which could partially explain why their win rates are higher.

D. Learning

The model of section III assumes that players observe their opponents' skill levels without error. If players instead gradually learn their opponents' skill, the econometric model of section IV is misspecified. In a learning model, the player updates his beliefs about his opponents' skill each hand based on his prior belief and a signal of skill. The updated belief is a weighted average of the prior and the new signal, and the weight assigned to the prior increases with the total number of signals observed. Thus, late signals shift the belief on average by less than early signals.

Specifically, a learning model implies that the effect of expected winnings decreases in magnitude as the history of signals becomes longer. I test for learning effects by estimating

$$(14)\quad \begin{aligned} M_{hti} &= \alpha_{ti} + \beta_1 \text{ActualWinnings}_{hti} + \beta_2 h_{hti} + \beta_3 h_{hti} \times \text{ActualWinnings}_{hti} \\ &+ \beta_3 X_{hti} + \varepsilon_{hti}. \end{aligned}$$

As in equation (10), h_{hti} is the number of hands played so far in session *t*. In the estimation of equation (14), $\text{OpponentSkill}_{hti}$ and $h_{hti} \times \text{OpponentSkill}_{hti}$ serve as instruments for the variables $\text{ActualWinnings}_{hti}$ and $h_{hti} \times \text{ActualWinnings}_{hti}$. The theoretical prediction of a learning model is that the negative effect of expected winnings should become weaker as a player learns more about his opponents, hence $\hat{\beta}_1 < 0$ and $\hat{\beta}_3 > 0$.

Table 6 shows the results for equation (14). Column 1 shows the OLS results. The main effect on dollars won is -0.002 and the effect between the interaction of dollars won and hundreds of hands played is 0.0004. Both effects are significantly different than zero. Column 2 shows the 2SLS results. The main effect on dollars won is -1.16 and the interaction effect is 0.67. While this pattern of coefficients is consistent with learning, the interaction effect is imprecisely estimated and is not statistically significant from zero. The no-learning model of section III cannot be rejected.

VII. Robustness Tests

A. Checking the Validity of the Instrument

For the results of the previous section to measure the effect of expected winnings on the probability of quitting, opponent VP$IP must be a valid instrument for expected winnings. The instrument is constructed by taking the average of each opponent's

VP$IP, but mean opponent VP$IP might not be a sufficient statistic for expected winnings. The first part of this section addresses this concern by allowing more flexibility in the first-stage relationship between opponent skill and actual winnings. Section III.B mentioned another threat to validity: VP$IP may affect the probability of quitting by changing the marginal utility of poker. The second part of this section addresses this concern by using an instrumental variable other than VP$IP.

1. Testing the Functional Form

Equation (11) specifies that actual winnings depend on mean opponent VP$IP and other controls. Since the instrument is constructed as the mean of several opponents' VP$IPs, I could instead use the individual VP$IP of each opponent as instruments. Such a specification could include the VP$IP of up to nine opponents, although this would limit the sample to games involving ten players. To keep a reasonably large sample size, I include the VP$IP of five opponents. That is, I estimate equation (10) by 2SLS using the first stage equation

$$\text{(15)} \quad \begin{aligned}\text{ActualWinnings}_{hti} = \delta_{ti} + \gamma_1 \text{OpponentVP\$IP}_1 + \ldots + \gamma_5 \text{OpponentVP\$IP}_5 \\ + \gamma_6 h_{hti} + \gamma_7 X_{hti} + u_{hti},\end{aligned}$$

where $\text{OpponentVP\$IP}_1$ is the VP$IP of the opponent to the player's immediate right, and $\text{OpponentVP\$IP}_5$ is the VP$IP of the fifth opponent to the player's right.

Table 7 shows the results. Column 1 shows the 2SLS, first-stage and reduced-form results when mean opponent VP$IP is used as the instrument for the sample of player-hands with at least five opponents. A one dollar increase in expected winnings decreases the probability of quitting by 0.94 percentage points. Column 2 shows the results when the VP$IPs of the five opponents to the player's right are used as instruments. A one dollar increase in expected winnings decreases the probability of quitting by 0.76 percentage points. The similarity of the estimates suggests that the effects estimated in early sections are not artifacts of the instrumental variable's functional form.

Since column 2 uses multiple instruments, I can check the specification by testing the overidentifying restrictions. The null hypothesis is that none of the instruments is correlated with the error term, and the test statistic is distributed chi-squared (n - 1), where n is the number of instruments[10]. The overidentification test statistic and its associated p-value are reported at the bottom of columns 2. The test suggests that the individual opponents' VP$IPs meet the requirements for a valid instrument outlined in section IV.

2. Using an Alternative Instrument

The overidentification test of table 7 might overlook correlation between individual opponents' VP$IPs and an omitted variable, however, because each opponent's VP$IP likely affects the same variables. To conduct a stronger test of VP$IP's validity as an instrument, I estimate the relationship between quitting and expected winnings using a different instrumental variable. Table 3 shows that high-experience players differ on many characteristics from low-experience players. High-experience players seem to be more skilled since they lose significantly fewer dollars per hand. Conditional on VP$IP, it may be difficult for players to distinguish between high- and low-experience opponents, so experience should be unlikely to affect the marginal utility of poker. The dummy Experience equals one for high-experience players. I define OpponentExperience to be the mean of Experience of a player's opponents. I then

[10] The reported test statistic is the Hansen-Sargan statistic, which measures the correlation between the instrumental variables and the residuals from the second stage.

estimate equation (10) using opponent experience as an instrument for expected winnings.

Table 8 reports the 2SLS, first-stage and reduced-form results. Column 1 reports the results when opponent VP$IP is omitted. A one dollar increase in expected winnings decreases the probability of quitting by 0.66 percentage points. Column 2 reports the results when the VP$IPs of five opponents are included as instruments in addition to opponent experience. A one dollar increase in expected winnings decreases the probability of quitting by 0.73 percentage points. Column 2 also reports the overidentification test statistic, which has a *p*-value of 0.16. These results suggest that that the results in section VI stem from variation in expected winnings and not in some omitted variable.

B. Checking the Functional Form of the Econometric Model

The main econometric model, equation (10), specifies that the probability of quitting follows a linear probability model. If instead some nonlinear function relates the regressors to the probability of quitting, the results from section VI could be misleading. This section estimates two nonlinear models: a probit model and a proportional hazards model.

1. Probit Results

Defining $\Phi(.)$ from equation (8) as the standard normal cumulative distribution function implies the following probit model:

$$(16) \quad \Pr[M_{hti} = 1] = \Phi(\alpha + \beta_1 \text{ActualWinnings}_{hti} + \beta_2 h_{hti} + \beta_3 X_{hti}),$$

where $\Phi(.)$ is the standard normal cdf. The large number of sessions makes it impractical to estimate session effects. I estimate equation (16) using standard probit techniques and also using a probit instrumental variables estimator, with opponent skill as the instrument. The probit IV technique is developed in Newey (1987).

Table 9 reports the probit results. The normalized probit estimate is $100^* \hat{\beta} \cdot \phi(\hat{\beta}\overline{X})$, where $\phi(.)$ is the standard normal density, which gives the marginal percentage point effect of X on the probability of quitting, evaluated at the sample mean of X. Column 1 shows the standard probit results. A one dollar increase in actual winnings decreases the probability of quitting by 0.0025 percentage points. Column 2 shows the probit IV results. A one dollar increase in expected winnings decreases the probability of quitting by 1.17 percentage points. For comparison, column 3 shows the linear probability 2SLS results when session effects are omitted. A one dollar increase in expected winnings decreases the probability of quitting by 1.41 percentage points. The similarity of the probit and linear probability model results bolsters the evidence from section VI.

2. Proportional Hazard Model Results

The linear probability and probit models specify that the regressors additively affect the probability of quitting. I estimate an alternative model in which the regressors multiplicitavely affect the probability of quitting:

(17) $\Pr[M_{hti} = 1] = \lambda_0 (h_{hti}) \exp(\beta X_{hti}),$

where $\lambda_0(h_{hti})$ is an unspecified baseline hazard function and X_{hti} is a vector of control variables[11]. A change in one of the control variables affects the probability of quitting by multiplying the baseline hazard.

[11] An "on the job" search model of poker, in which the poker player's outside option varies over time and the player stops playing when the outside option exceeds a certain threshold, could provide theoretical justification for a proportional hazard model. See section 4.3 of Van Den Berg (2001).

Table 10 reports the proportional hazard results. The reported estimates are hazard ratios, which give the multiplicative effect on the baseline probability of quitting of a one unit change in the control variable. The hazard ratio equals $\exp(\hat{\beta})$. Column 1 shows the results using actual winnings and a set of dummies for number of players as controls. A one dollar increase in actual winnings multiplies the baseline hazard by 0.999. For a baseline hazard of 2.4 percent (the sample average), this implies a 0.0024 percentage point decrease in the probability of quitting. Column 2 shows the results using opponent VP$IP*20 and number of player dummies as controls. A five percentage point increase in opponent VP$IP multiplies the baseline hazard by 0.782. For a baseline hazard of 2.4 percent, this implies a 0.52 percentage point decrease in the probability of quitting. For comparison, column 4 of table 2 implies that a five percentage point increase in opponent VP$IP decreases the probability of quitting by 0.95 percentage points. Compared to the linear probability model, the proportional hazard model gives a significantly smaller estimate of the reduced form relationship between opponent VP$IP and the probability of quitting.

VIII. Conclusion

This paper finds that the skill of a poker player's opponents predicts the amount the player will win. Players are more likely to quit playing poker against highly skilled opponents and thus more likely to quit playing when expected winnings are low. Using opponent skill as an instrumental variable for expected winnings, I estimate that a one dollar increase in expected winnings decreases the probability of quitting by 0.9 percentage points. This is a large effect relative to the mean probability of quitting of 2.4 percent, and it corresponds to a price elasticity of session length (the number of hands played at a given table) of -0.11. Substitution between poker tables, rather than between poker and other activities, accounts for roughly half of this effect, but expected winnings still affects the total amount of poker played. In fact, I calculate that the price elasticity of the total amount of poker played is -0.14, which is similar to the session-length elasticity.

The negative relationship between expected poker winnings and the amount of poker played has several broader implications. First, it shows that people adjust their leisure consumption in response to changes in the prices of leisure activities. Because this paper does not measure the total amount of leisure consumed or the labor supply of poker players, I cannot test whether changes in expected poker winnings change the total amount of leisure consumed, and thus labor supply, or merely cause players to reallocate time between poker and some other leisure activity, keeping total leisure constant. But the results do suggest a path by which the price of leisure could affect labor supply, which deserves further research.

Second, the results are consistent with rational poker players who realize the tradeoff between poker and consumption. When poker becomes too expensive, they quit playing. This contributes to the many studies finding that wagering markets predict outcomes of sporting events well (Sauer 1998) and that demand for lottery tickets increases in the expected value of lottery tickets (Kearny 2006). Together this evidence suggests that standard economic models describe at least some aspects of gambling. My paper also shows, however, that gambling behavior still presents many puzzles. The players in the data set employ a wide range of strategies, some of which perform much worse than others. Whether traditional economic models offer parsimonious and plausible explanations for poor poker play remains an open question.

References

AGA. (2006). "State of the States: The AGA Survey of Casino Entertainment," American Gaming Association, Washington, D.C.

Aguiar, Mark and Erik Hurst. (2006). "Measuring Trends in Leisure: The Allocation of Time over Five Decades," Federal Reserve Bank of Boston, Working Paper No. 06-2.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. (1997). "Labor supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics*, 112 (2): 407-41.

Chou, Yuan K. (2000). "Testing Alternative Models of Labor Supply: Evidence from Taxi Drivers in Singapore," University of Melbourne, Research Paper No. 768.

Clotfelter, Charles T. and Philip J. Cook. (1993). "The Peculiar Scale of Economies of Lotto," *American Economic Review*, 83 (3): 634-643

Connolly, Marie. (2006). "Here Comes The Rain Again: Weather and the Intertemporal Substitution of Leisure," Princeton University, Working Paper.

Farber, Henry. (2005). "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy*, 133 (1): 46-82.

Fehr, Ernst, and Lorenz Götte. (2002). "Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment," Institute for Empirical Research in Economics, University of Zürich, Working Paper No. 125.

Guryan, Jonatha and Melissa S. Kearney. (2005). "Lucky Stores, Gambling, and Addiction: Empirical Evidence from State Lottery Sales," NBER Working Paper No. 11287.

Kearney, Melissa S. (2006). "State Lotteries and Consumer Behavior," forthcoming, *Journal of Public Economics*.

List, John A. (2004). "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," *Quarterly Journal of Economics*, 119 (1): 49-89.

List, John A. (2003). "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics*, 118 (1): 41-71.

Neidell, Matthew. (2006). "Public Information, Avoidance Behavior, and Health: Do People Respond to Smog Alerts?" Columbia University, Working Paper.

Newey, Whitney. (1987). "Efficient Estimation of Limited Dependent Variable Models with Endogenous Explanatory Variables," *Journal of Econometrics*, 36: 231-250.

Oettinger, Gerald S. (1999). "An Emprical Analysis of the Daily Labor Supply of Stadium Vendors," *Journal of Political Economy*, 107 (2): 360-92.

Sauer, Raymond D. (1998). "The Economics of Wagering Markets," *Journal of Economic Literature*, 36 (4): 2021-2064.

Van De Berg, Gerard J. (2001). "Duration Models: Specification, Identification and Multiple Durations," in, J. J. Heckman and E. Leamer, eds., *Handbook of Econometrics*, Vol. 5 (North-Holland, Amsterdam) pp. 3382-3460.

Xanthopoulos, Judy. (2006). "Internet Poker Industry and Revenue Analysis: Final Report," Poker Players Alliance, Working Paper.

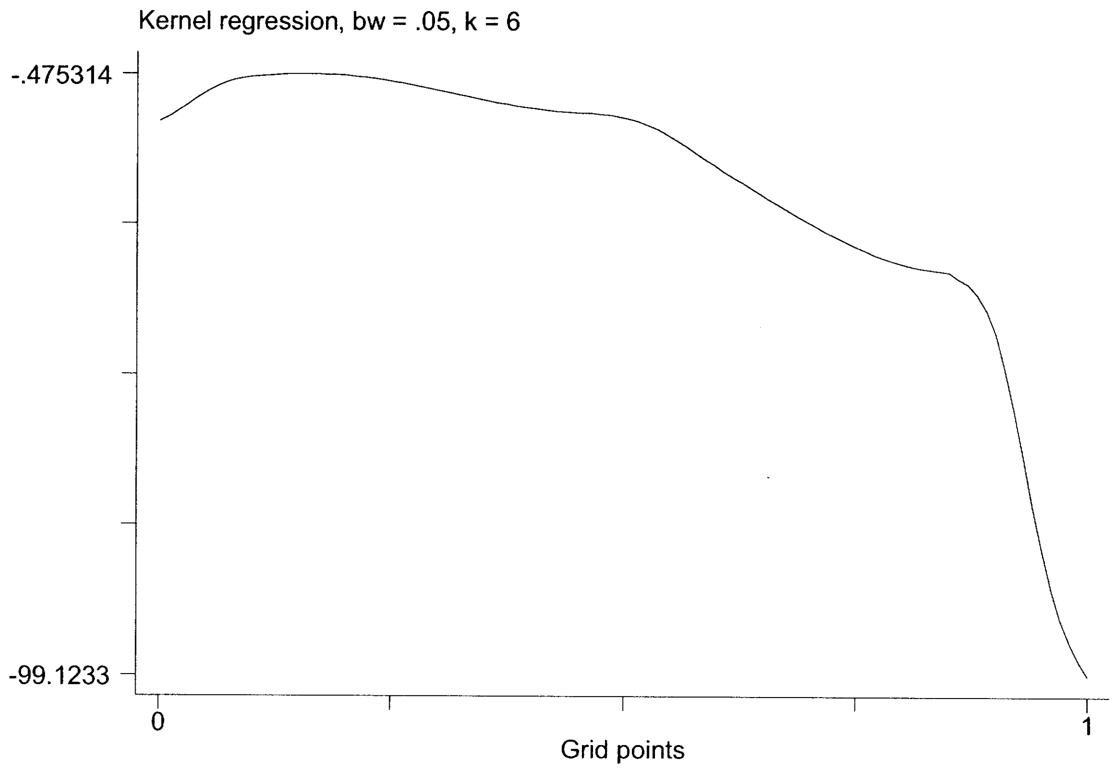

FIGURE 1
Kernel Regression of Dollars Won per Hand on Own VP$IP

TABLE 1
Descriptive Statistics

	Mean (1)	Median (2)	Std Dev (3)
Panel A: Player summary statistics (N=4,586)			
Total sessions	12.2	3	28.1
Total hands played	519.3	130	1207.5
Minutes of playing time	390.7	93.6	920.7
Average session length	31.1	23.7	29.4
Total dollars won	-148.3	-300	2593.7
Dollars won per hour of play	-571.2	-173.4	2485.0
Dollars won per hand played	-7.2	-2.1	29.7
VP$IP	0.37	0.34	0.17
Panel B: Session summary statistics (N=55,983)			
Session length in hands	42.5	26	50.9
Session length in minutes	32.0	18.9	39.0
Total dollars won	-12.15	-30	702.30
Panel C: Player-hand summary statistics (N=2,381,396)			
Last hand	0.024	0	0.152
Dollars won	-0.3	0	107.7
Opponent VP$IP	0.276	0.266	0.051
Number of players at table	8.1	9	1.9
Cumulative hands in current session	52.2	31	64.7
Cumulative dollars won in current session	20.2	-30	747.6
Cumulative session dollars per hand	0.012	-1.243	36.130

TABLE 2

Relationship between Probability of Quitting and Expected Winnings, Instrumented with Opponent VP$IP

	Dependent Variables				
	Dollars Won in Current Hand: OLS (1st stage)	Last Hand of Session: OLS	Last Hand of Session: OLS	Last Hand of Session: OLS (Reduced Form)	Last Hand of Session: 2SLS
	(1)	(2)	(3)	(4)	(5)
Opponent VP$IP	21.094			-19.01	
	(2.566)			(.4616)	
Dollars won in current hand		-0.0021			-0.9013
		(.0001)			(.1106)
Dollars won up to current hand per hand			-0.0024		
			(.0004)		
Hands played	-0.0095	0.0275	0.0276	0.0259	0.0174
	(.0018)	(.0015)	(.0015)	(.0014)	(.0022)
Observations	2,363,348	2,363,348	2,363,348	2,363,348	2,363,348
Number of player effects?	Yes	Yes	Yes	Yes	Yes
Session effects?	Yes	Yes	Yes	Yes	Yes

Notes: Coefficients are predicted percentage point change in the probability that the current hand is the last hand of a player-session. Robust standard errors in parentheses account for clustering at the session level.

TABLE 3
Relationship between Probability of Logging Off and Expected Winnings

	Full Sample		Restricted Sample			
	Logs Off: OLS	Logs Off: 2SLS	Logs Off: OLS	Logs Off: 2SLS	Last Hand of Session: OLS	Last Hand of Session: 2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Dollars won in current hand	-0.0016	-0.4729	-0.0018	-0.5434	-0.0020	-1.0268
	(.0001)	(.0585)	(.0003)	(.1885)	(.0005)	(.3520)
Hands played	0.0151	0.0098	0.0284	0.0209	0.0400	0.0350
	(.0008)	(.0012)	(.0046)	(.0054)	(.0070)	(.0096)
Observations	2,363,348	2,363,348	120,726	120,726	120,726	120,726
Number of player effect?	Yes	Yes	Yes	Yes	Yes	Yes
Session effects?	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Coefficients are predicted percentage point change in the probability that the current hand is the last hand of a player-session. Robust standard errors in parentheses account for clustering at the session level.

TABLE 4

Differences between High- and Low-Experience Players

Experience	All	High	Low
	(1)	(2)	(3)
Total dollars won	-148	502	-383
	(2594)	(4258)	(1554)
Total hands played	519	1495	168
	(1207)	(1998)	(280)
Total sessions	12.2	34.4	4.2
	(28.1)	(47.2)	(5.8)
Minutes of playing time	391	1124	126
	(921)	(1530)	(216)
Average session length	31.1	34.2	29.9
	(29.4)	(21.2)	(31.8)
Dollars won per hour	-571	-16	-771
	(2485)	(521)	(2855)
Dollars won per hand	-7.25	-0.19	-9.79
	(29.71)	(6.06)	(34.11)
VP$IP	0.37	0.28	0.40
	(0.17)	(0.10)	(0.18)
Total players	4586	1215	3371

Notes: Players are classified as experienced if they are ever observed playing at multiple tables simultaneously. Standard deviations in parentheses.

TABLE 5
Relationship between Probability of Quitting and Expected Winnings by Experience

	Dependent Variables				
	Dollars Won in Current Hand: OLS (1st stage)	Last Hand of Session: OLS	Last Hand of Session: 2SLS	Logs Off: OLS	Logs Off: 2SLS
	(1)	(2)	(3)	(4)	(5)
Panel A: Low-Experience Players					
Opponent VP$IP	43.065				
	(5.814)				
Dollars won in current hand		-0.0040	-0.4281	-0.0035	-0.3571
		(.0001)	(.0579)	(.0001)	(.0488)
Hands played	-0.0380	0.0333	0.0151	0.0289	0.0138
	(.0055)	(.0021)	(.0031)	(.0018)	(.0026)
Observations	560,108	560,108	560,108	560,108	560,108
Number of player effects?	Yes	Yes	Yes	Yes	Yes
Session effects?	Yes	Yes	Yes	Yes	Yes
Panel B: High-Experience Players					
Opponent VP$IP	13.522				
	(2.838)				
Dollars won in current hand		-0.0012	-1.4236	-0.0007	-0.5931
		(.0001)	(.3008)	(.0001)	(.1261)
Hands played	-0.0024	0.0261	0.0211	0.0117	0.0096
	(.0018)	(.0017)	(.0030)	(.0008)	(.0013)
Observations	1,803,240	1,803,240	1,803,240	1,803,240	1,803,240
Number of player effects?	Yes	Yes	Yes	Yes	Yes
Session effects?	Yes	Yes	Yes	Yes	Yes

Notes: Coefficients are predicted percentage point change in the probability that the current hand is the last hand of a player-session. Robust standard errors in parentheses account for clustering at the session level.

TABLE 6

The Interaction between Expected Winnings and Hands Played

	Dependent Variables	
	Last Hand of Session: OLS	Last Hand of Session: 2SLS
	(1)	(2)
Dollars won in current hand	-0.0023	-1.1567
	(.0001)	(.1719)
Hands played	0.0275	0.0270
	(.0015)	(.0071)
Dollars won in current hand * hundreds of hands played	0.0004	0.6682
	(.0001)	(.4033)
Observations	2,363,348	2,363,348
Number of player effects?	Yes	Yes
Session effects?	Yes	Yes

Notes: Coefficients are predicted percentage point change in the probability that the current hand is the last hand of a player-session. Robust standard errors in parentheses account for clustering at the session level.

TABLE 7
Instrumenting for Expected Winnings
with VP$IPs of Several Opponents

	(1)	(2)
Panel A: 2SLS relationship between expected winnings and quitting		
Dollars won in current hand	-0.937	-0.757
	(.108)	(.086)
Panel B: First stage		
Mean opponent VP$IP	27.53	
	(3.15)	
Opponent 1 VP$IP		3.16
		(1.21)
Opponent 2 VP$IP		4.86
		(1.08)
Opponent 3 VP$IP		5.75
		(1.03)
Opponent 4 VP$IP		5.58
		(.98)
Opponent 5 VP$IP		3.52
		(.94)
Panel C: Reduced-form		
Mean opponent VP$IP	-25.80	
	(.54)	
Opponent 1 VP$IP		-4.34
		(.18)
Opponent 2 VP$IP		-3.93
		(.17)
Opponent 3 VP$IP		-3.58
		(.15)
Opponent 4 VP$IP		-3.70
		(.15)
Opponent 5 VP$IP		-2.91
		(.14)
Overidentification test statistic		6.09
p-value		0.19

Notes: N=2,068,544. All regressions control for hands played, number of player effects and session effects. Coefficients are predicted percentage point change in the probability that the current hand is the last hand of a player-session. Robust standard errors in parentheses account for clustering at the session level.

TABLE 8

Overidentification Test with Opponent Experience

	(1)	(2)
Panel A: 2SLS relationship between expected winnings and quitting		
Dollars won in current hand	-0.660	-0.726
	(.100)	(.077)
Panel B: First stage		
Mean opponent experience	-5.62	-3.12
	(.83)	(.90)
Opponent 1 VP$IP		2.32
		(1.24)
Opponent 2 VP$IP		4.03
		(1.10)
Opponent 3 VP$IP		4.96
		(1.05)
Opponent 4 VP$IP		4.86
		(1.00)
Opponent 5 VP$IP		2.93
		(.96)
Panel C: Reduced form		
Mean opponent experience	3.71	1.54
	(.14)	(.15)
Opponent 1 VP$IP		-3.93
		(.18)
Opponent 2 VP$IP		-3.52
		(.17)
Opponent 3 VP$IP		-3.19
		(.15)
Opponent 4 VP$IP		-3.34
		(.15)
Opponent 5 VP$IP		-2.62
		(.14)
Overidentification test statistic		7.90
p-value		0.16

Notes: N=2,068,544. All regressions control for hands played, number of player effects and session effects. Coefficients are predicted percentage point change in the probability that the current hand is the last hand of a player-session. Robust standard errors in parentheses account for clustering at the session level.

TABLE 9
Probit Estimates of the Relationship between Probability of Quitting and Expected Winninngs

	Dependent Variables		
	Last Hand of Session: Probit	Last Hand of Session: Probit IV	Last Hand of Session: LPM 2SLS
	(1)	(2)	(3)
Dollars won in current hand	-0.0025	-1.1703	-1.4128
	(.0001)	(.0239)	(.2866)
Hands played	-0.0062	-0.0074	-0.0062
	(.0002)	(.0002)	(.0016)
Observations	2,363,348	2,363,348	2,363,348
Number of player effects?	Yes	Yes	Yes

Notes: Probit estimates are normalized to reflect the marginal percentage point effect at X^* of X on the probability of quitting. The normalized probit estimate is $100 \times \hat{\beta} \times \phi(X * \hat{\beta})$, where $\phi(\cdot)$ is the standard normal density.

TABLE 10

Proportional Hazard Model of Ending Sessions

	Hazard Ratios	
	(1)	(2)
Dollars won in current hand	0.999	
	(.000)	
Opponent VP$IP*20		0.782
		(.004)
Observations	2,363,348	2,363,348
Number of player effects?	Yes	Yes

Notes: Estimates are hazard ratios, the exponetiated coefficients from the proportional hazards model. Robust standard errors in parentheses account for clusering at the player-session level.

Undeclared: Patterns of Major Choice in College

Timothy M. Watts

Abstract

This paper develops and tests a model of how college students choose their field of study. The model combines features from learning and human capital models and captures several stylized facts from the empirical literature on choice of college major. I test the model's predictions using High School and Beyond data. I find three results that generally agree with the model's predictions. First, students with higher levels of ability choose majors with higher average earnings. Second, students who receive low grades in college are more likely to change their field of study. Third, students who switch majors in college subsequently earn less than students who do not change majors, but this difference is primarily due to major-switchers obtaining degrees in low-paying fields.

I. Introduction

In this chapter I develop a model that focuses on the role of learning and occupation-specific human capital in the choice of college major. The model assumes that there are two occupations with different returns to ability. Students have either high or low ability, but do not know their ability with certainty. Based on their beliefs about their ability, students choose an initial field of study for college. Studying a field in college generates occupation-specific human capital that increases a student's potential earnings in that field's associated occupation. Midway through college, students learn their ability perfectly and may then switch their field of study. Upon the completion of college, students enter the labor force and choose the occupation that gives them the higher earnings.

The model is consistent with several well documented results and generates several predictions about the pattern of major choice in college. The model predicts that students with high expected ability begin their college studies in the field with higher returns to ability; that the number of students who switch from the high ability field to the low ability field exceeds the number of students who switch in the opposite direction; that students who receive low grades in college are likely to switch their field of study; and that students who switch their field of study earn less than those who do not.

I test these predictions using High School and Beyond data. I find that students who have higher scores on a high-school test of cognitive ability plan to major in fields with higher average earnings. I also find that students who receive low grades in college are more likely to change their major field of study. Finally, I find that students who switch majors in college earn less than students who do not, although this result is explained by differences in the distribution of major choices for switchers and non switchers. On the whole, these results support the model.

II. Evidence on Field of Study in College

Prior research on the choice of majors in college documents several robust empirical findings. First, average earnings vary greatly across majors, and this result persists across different datasets and different cohorts of college students. To give a few examples, differences in average earnings across college majors exist for American college students who graduated from college in the year 1986 (Rumberger and Thomas, 1993); Canadian students who graduated in the years 1982, 1986 and 1990 (Finnie and Frenette, 2003); and students who graduated from the University of Texas in five different years spanning the time period 1980 to 2000 (Hamermesh and Donald, 2006). These earnings differences persist even after controlling for occupation, college performance, achievement test scores and a wide range of demographic variables. These studies also consistently identify particular majors as having high or low earnings. For example, students who major in the fields of engineering or health earn above average wages, while those who major in the fields of humanities or education earn below average wages.

The literature also finds that differences in average earnings across majors influence a student's choice of major. Freeman (1971) finds that the number of college students who major in the field of engineering rises with the market wages of engineers. Subsequent studies find similar results for students studying law (Freeman, 1975), economics (Hoffman and Low, 1983) and education (Zarkin, 1985). Berger (1988) considers the general problem of major choice and estimates separate wage equations for several college fields of study. He finds that a student's probability of choosing a given major increases with the student's expected present value of income in that major.

Prior research also provides some evidence that students sort themselves across majors by ability. Paglin and Rufolo (1990) find that mean math and verbal scores on the Graduate Record Exam (GRE) vary widely across majors. They also find a strong positive correlation between major fields' mean starting salaries and mean GRE math scores. The second result suggests that students with higher math ability choose majors with higher average salaries, although the authors acknowledge that a student's choice of major could influence her math ability since some majors could increase math ability more than others and the GRE only measures a student's ability at the end of college. The possibility that employers place a high value on math ability relative to other skills provides another possible explanation for the correlation between mean major salaries and mean major math scores.

A related strand of literature examines students who switch majors in college. Two large national surveys conducted in the 1960s show that many students change their career plans in college. The first survey finds that 37% of college students change their career plans between their freshman and senior year (Davis, 1965). The second survey finds that 38% of male college students who choose certain fields of study during their freshman year change their career plans by their sophomore year (Werts, 1967)[1]. To the extent that students' career plans correspond to their choice of majors, both surveys

---

[1] This statistic is based on my calculations using data reported in Table 1 of Werts (1967). The calculation likely overestimates the proportion of students who change career plans because the table only reports data for careers fields in which at least 150 of those initially interested in the field subsequently switch into another field.

suggest that students often begin college in one field of study only to graduate with a degree in a different field of study.

When students change their career plans, they choose some career fields more frequently than others. This means that some fields grow while others shrink. Davis (1965) finds that, between the freshman and senior year of college, the fields of engineering and medicine lose the most students on net, while the fields of education and business gain the most students on net. Net changes in field size, however, do not fully describe the pattern of career choice. Each field loses some of the students who initially choose it and gain some students who initially choose a different field.

In summary, the literature on the choice of majors broadly supports the finding that students who major in some fields earn significantly more than students majoring in other fields, even after controlling for observable characteristics. The variation in earnings across major fields partly explains a student's choice of majors, and students tend to choose a major in which their expected earnings are relatively high. Students' abilities also partly explain their choice of major, and there is weak evidence that the ranking of fields by average student ability corresponds to the ranking of fields by average expected earnings. Finally, a large proportion of students change their field of study in college, and some fields gain students on net while other fields lose students on net.

III. A Theory of College Major Choice

This section adapts Gibbons and Waldman's (1999) model of wage dynamics inside firms to the context of choosing a course of study in college. The present model differs from Gibbons and Waldman's model in two respects. First, the present model incorporates career-specific human capital to capture the idea that taking courses in one major likely increases a student's productivity in some occupations more than others. Second, in order to simplify the analysis, this model incorporates perfect rather than gradual learning.

A. The Model

The model lasts for three periods. In the first period, students take college courses in one field of study and receive grades. At the end of the first period, students decide whether to continue in their initial field of study or to switch fields. In the second period, students finish college and they enter the labor market in the third period.

Students differ in their innate ability, θ, which affects their productivity in both careers. For simplicity, innate ability takes either a high (θ_H) or low (θ_L) value. Before beginning college, students do not know their ability with certainty but they have a prior belief p_0, where $p_0 = \Pr[\theta = \theta_H]$. A student may derive her prior belief from a variety of sources, including academic performance, achievement test scores and family background characteristics. Prior beliefs are informative about actual ability in the following sense: of the population of students with prior belief p_0, the fraction p_0 has high ability and the fraction $(1 - p_0)$ has low ability. The grades that a student receives at the end of the first period perfectly inform her about her ability. I assume that half of all students have high ability and that the population distribution of p_0 is $U[0,1]$.

Students choose between two possible fields of study in college, called field 1 and field 2. Each field is associated with an occupation in the labor market. Students accumulate one unit of human capital specific to a field of study for each period they spend at college in that field.

When students enter the labor market, they work in occupation 1 or occupation 2 and earn an amount equal to their marginal revenue product. A student's productivity in either occupation depends both on the student's innate ability and the student's stock of human capital in the field of study associated with that occupation. Specifically, earnings in career j for student i with ability θ_i and occupation-specific human capital x_{ij} equal

$$y_{ij} = c_j + d_j(\theta_i f(x_{ij})) \quad (1)$$

where $x_{ij} \in \{0,1,2\}$ and $f(2) > f(1) > f(0)$. I assume that $d_2 > d_1 \geq 0$, that $c_1 > c_2$ and that $\theta_H > \frac{c_1 - c_2}{d_2 - d_1} > \theta_L$. Together these assumptions ensure that, among students who have equal stocks of human capital in both occupations, those with low ability earn more in occupation 1 and those with high ability earn more in occupation 2.

In each period, students make decisions that maximize the expected value of their lifetime earnings. In period 3, students know their ability and stock of human capital with certainty, and they choose the occupation that gives them higher earnings.

In the second period, students know their ability with certainty but must decide whether to accumulate a second unit of human capital in their initial field of study or to accumulate one unit of human capital in a new field. Students compare their earnings conditional on having two units of human capital in their initial field with their earnings conditional on having one unit of human capital in each field (both conditional on their known level of ability). Here I assume that $\theta_H > \frac{c_1 - c_2}{d_2 f(1) - d_1 f(2)}$

and $\theta_L > \frac{c_1 - c_2}{d_2 f(2) - d_1 f(1)}$, which ensures that in the second period high-ability students choose to study field 2 and low-ability students choose to study field 1.

In period 1, students calculate their expected third-period earnings for both initial fields of study. The expected earnings for a student with prior p_0 who begins her studies in field 1 is

(2) $$V_1 = (1 - p_0)[c_1 + d_1\theta_L f(2)] + p_0[c_2 + d_2\theta_H f(1)].$$

The first term of the equation is the probability that a student learns she has low ability multiplied by the expected earnings of a low ability student who initially studies field 1. The second term is the probability that the student learns she has high ability multiplied by the expected earnings of a high ability student who initially studies field 1. By the same reasoning, expected earnings for starting in field 2 equals

(3) $$V_2 = (1 - p_0)[c_1 + d_1\theta_L f(1)] + p_0[c_2 + d_2\theta_H f(2)].$$

Students choose their initial field of study to maximize expected earnings, so students only begin their studies in the second field if $V_2 > V_1$. Subtracting V_1 from V_2 and simplifying yields

(4) $$V_2 - V_1 = [p_0(d_2\theta_H + d_1\theta_L) - d_1\theta_L](f(2) - f(1)),$$

which is clearly increasing in p_0. Denote the prior at which $V_2 = V_1$ as p^*. Then students with $p_0 \le p^*$ begin their studies in field 1 and those with $p_0 > p^*$ begin their studies in field 2. Equation (4) and the definition of p^* imply that

(5) $$p^* = \frac{d_1\theta_L}{d_2\theta_H + d_1\theta_L} < \frac{1}{2},$$

where the inequality follows from the previous assumptions that $d_2 > d_1$ and $\theta_H > \theta_L$.

B. Predictions and Discussion

This simple model makes several predictions about the pattern of major choice and the relationship between major choice and earnings. First, students with high expected ability in high school will begin their studies in a field with high returns to ability.

Second, more students will switch from high-ability majors to low-ability majors than vice versa. This prediction implies that high-ability majors lose students on net. To see why this is true, note that any student with $p_0 > p^*$ studies the high-ability field in the first period. Since $p_0 \sim U[0,1]$ and $p^* < 1/2$, more than half of the students begin in the high ability field. But only half of the students—those who learn that they have high ability—study the high-ability field in the second period. Thus fewer students study the high-ability field in the second period than in the first.

Third, a student's college grades will be negatively correlated with the probability of switching majors. In the model, students perfectly learn their ability in college. If a student learns their ability through college grades, then receiving high grades reveals that the student has high ability and receiving low grades reveals that the student has low ability. Both high grades and low grades may cause a student to switch majors. But because the high-ability major loses students on net and because low grades cause students to switch out of the high-ability major, most student who switch majors will have low grades.

Fourth, students who switch majors earn less than those who do not. Two factors contribute to this result. First, major switchers necessarily accumulate less human capital in their chosen occupations than do non-switchers. This difference implies that, within occupations, major switchers earn less than non-switchers (because earnings in both occupations increase in occupation-specific human capital). Second, the majority of major switchers end up working in the low-ability, low-earnings occupation, which further deflates the average earnings of major switchers.

IV. Data

This paper uses the High School and Beyond dataset (HSB). HSB participants were either sophomores or seniors in high school in 1980, the base year of the survey. HSB conducted follow-up surveys for both cohorts in 1982, 1984 and 1986 and for the sophomore cohort in 1992. Because there is limited information on labor market outcomes for the senior cohort, I limit my sample to the sophomore cohort. I further restrict my sample to students who earn a bachelor's degree by the end of the sample period and have complete information on planned and actual college major. The main sample consists of 2,384 students. For the analysis of labor market outcomes, I further restrict my sample to students who make their first long-term transition to the labor force during the sample period. Following Farber and Gibbons (1996), I define a long-term transition to the labor force as spending at least three consecutive years working after spending at least one year not working. The sample for labor market outcomes comprises 1,775 students. The appendix provides more detail on the construction of the sample.

To examine major switching behavior, I define variables for college students' planned and actual majors. The 1982 HSB follow-up survey asks members of the sophomore cohort to choose the field that they "would most like to study in college" from a list of 25 fields. I define a student's planned major as her response to this question. HSB also contains transcript data from postsecondary education institutions attended by participants. HSB records the three-digit Classification of Instructional Program (CIP) codes for the major field of study listed on each postsecondary transcript. To determine actual college major, I first assign each CIP code to the one of the 25 fields for planned college major[2]. I then define actual college major as the field corresponding to the CIP code from the earliest transcript that records a completed bachelor's degree. If a student's actual major differs from her planned major, I define her as having switched majors.

In addition to the data on students' major field of study, I use demographic and labor market variables. The demographic variables include gender and race as well as two measures of ability: score on a standardized cognitive test and college GPA. The cognitive test score is a composite of reading, vocabulary and math scores on an ability test administered by HSB in 1982. A student's college GPA is taken from the same transcript as actual college major and converted to a 4.0 scale. The timing of the ability measures plays an important role in the analysis. The cognitive test score is determined before students make their initial major choices, while the college GPA is determined (at least in part) between the students' initial and final major choices[3].

Some regressions also control for characteristics of the student's family and high school. Family background variables include indicators for whether the student's father or mother has a college degree, whether the student's family's income was greater than $25,000 in 1980, and whether a language other than English was spoken at the student's

---

[2] The appendix gives more information on the assignment of CIP codes to major fields.

[3] Since the final major choice occurs before the end of college, then part of college GPA is determined after the final choice. Even so, it contains information about ability that is unknown at the time of initial major choice.

home. High school variables include indicators for whether the school is in the southern Census region, whether the school is private and whether the school is in an urban area.

The analysis of labor market outcomes focuses on average monthly earnings. I construct average monthly earnings using data from years that satisfy two criteria: they must occur after an individual's transition to the labor force and they must be spent mainly working[4]. I refer to these years as working years. For each working year from 1986 to 1992, I determine how many months an individual spends working full time using monthly data on labor force status and educational enrollment status[5]. I then calculate average monthly earnings by dividing annual earnings by the number of months spent working full time. The data contain up to seven years of average monthly earnings for each worker. I control for an individual's experience and age in earnings regressions. I define experience as the number of years that have elapsed since a worker's transition to the labor market. Because students complete the same level of education at different ages, I can separate the effects of age and experience on labor market outcomes.

A. Student Characteristics by College Major

Table 1 reports average characteristics of students by planned and actual major. Panel A shows that business is the most common planned major, accounting for 21% of the students in the sample, followed by engineering (12%), preprofessional (11%) and computer science (8%). Majors also differ widely in gender balance, ranging from predominantly male fields such as engineering to predominantly female fields such as

---

[4] Individuals must work full time for at least four months in 1992 and at least six months in all other years to be classified as mainly working for that year.

[5] Labor force status is available from January 1982 to August 1992 and education enrollment status is available from June 1982 to June 1992.

home economics. The proportion of black students also varies across majors, although not as widely.

Panel B shows characteristics of students by actual majors. Business remains the most popular major. In fact, it has grown to include 25% of students. The social sciences (13%) and engineering (8%) follow business in popularity. The gender and racial makeup of actual majors resembles that of planned majors.

B. Pattern of Major Switching

Table 2 reports the difference between the number of students who plan to major in a field and the number of students who actually receive a degree in that field. The social sciences, business and English gain the most students, while preprofessional, engineering and computer sciences lose the most students. As a percentage of the number of students who plan to major in a field, the social sciences, English and home economics experience the biggest gains, while preprofessional, vocational and health occupations experience the biggest losses. The breakup of fields into net gainers and losers broadly resembles that found in earlier studies (Davis, 1965, and Werts, 1967).

Table 2 also shows that the preprofessional field suffers remarkable attrition. Over 250 students plan to major in the field but only 5 actually do. HSB's survey question for planned major lists prelaw, premedicine and predentistry as examples of preprofessional fields. Few schools award majors in these fields, however, and aspiring doctors commonly major in fields such as biology. Because many of these students are likely to be wrongly classified as major switchers, I exclude students who plan to major in the preprofessional field from subsequent analysis.

C. Major Switchers versus Non-Switchers

Table 3 compares the characteristics of major switchers and non-switchers. Planned major predicts actual major fairly well – 44% of students do not switch majors. The proportion of females is six percentage points higher in switchers than in non-switchers, but the other background characteristics are relatively similar for the two groups. Switchers are somewhat less likely to make a long-term transition to the labor force.

V. Empirical Results

This section reports how the model's predictions stand up to the data. I first investigate the relationship between a field of study's mean student ability and mean earnings. The model predicts that students with high expected ability begin their college studies in fields with high returns to ability. The prediction also implies that students with high expected ability begin their college studies in fields with high average pay[6].

I test this prediction by regressing the expected earnings of a student's planned major on the student's score on the cognitive test administered during high school. I define a major's expected earnings as the average monthly salary offered during July of 1982 to students graduating with a bachelor's degree in that major. Table 4 shows the results. Students with higher test scores plan to major in fields with higher starting salaries. A student whose test score is one standard deviation above the mean selects a major with a monthly starting salary that is $86 higher than a student whose test score is one standard deviation below the mean. When I add controls for the student's demographic characteristics, the effect of ability on the planned major's starting salary

[6] Let $y_j(\theta_i, x_1, x_2)$ denote the earnings in occupation j for a student with ability θ_i and x_1 units of human capital specific to occupation 1 and x_2 units of human capital specific to occupation 2. By assumption, $y_2(\theta_H,1,1) > y_1(\theta_H,2,0)$. Because returns to ability are positive in both occupations, $y_1(\theta_H,2,0) > y_1(\theta_L,2,0)$. And because returns to specific human capital are positive in both occupations, $y_2(\theta_H,0,2) > y_2(\theta_H,1,1)$ and $y_1(\theta_L,2,0) > y_1(\theta_L,1,1)$. Together, these inequalities give the following ranking of earnings for the four ability-capital combinations , $y_2(\theta_H,0,2) > y_2(\theta_H,1,1) > y_1(\theta_L,2,0) > y_1(\theta_L,1,1)$. Thus all students who enter occupation 2 earn more than all students who enter occupation 1 and average earnings in occupation 2 exceed average earnings in occupation 1.

becomes somewhat smaller but remains highly statistically significant. These results support Paglin and Rufolo's (1990) finding that students with higher GRE math scores major in fields with higher average earnings, and my data overcome two of the weaknesses of their analysis. First, my measure of ability is observed before students enter college, so college curriculum does not cause the variation in ability. Second, I measure the majors' average starting salaries during the students' final year of high school. This ensures that the variation in the students' skill across planned major fields does not directly influence the fields' salaries because the students are not yet in the labor force.

The model also predicts that receiving low grades in college increases the probability that a student switches majors. I test this prediction by regressing an indicator for switching majors on college GPA and other controls. Table 5 reports the results. In the bivariate regression, college GPA has a small negative but statistically insignificant effect on the probability of switching majors. When I control for demographic characteristics and a set of effects for planned major field of study, the magnitude of the effect increases. Increasing college GPA by one point decreases the probability of switching majors by 6.8 percentage points, and this effect is significant at the 95% confidence level. Adding cognitive test score to the set of controls decreases the effect of college GPA, but it remains negative and statistically significant at the 90% confidence level. Taken together, the results provide some evidence that students switch majors when they receive low grades in their planned major.

I next examine the relationship between switching majors and earnings. The model predicts that major switchers earn less than non-switchers even after controlling for actual major field. I estimate the following wage equation:

(6) $I_{it} = a_t + \delta_1 S_i + \delta_2 X_{it} + \mu_{it}$

The dependent variable I_{it} is the logarithm of average monthly earnings for individual i in year t. a_t is a year fixed effect, S_i is an indicator for switching majors, X_{it} is a vector of other control variables and μ_{it} is the error term.

The first column of table 6 shows that major switchers earn 7.5 percent less than non-switchers after controlling for experience, age, gender, race and year effects. The estimate remains nearly unchanged in column 2, which includes measures of student ability as additional independent variables. When a set of effects for actual college major is included (column 3), the negative effect of switching majors on earnings drops to 2.2 percent and becomes statistically insignificant. When I include industry or occupation effects instead of actual major effects (columns 4 and 5), I estimate the negative impact switching majors on earnings to be 6.4 or 5.6 percent. Finally, including a full set of effects for actual major, industry and occupation (column 6) yields a negative switching effect of 2.2 percent, which is not statistically different from zero.

These results provide mixed support for the model. In each specification, switching majors has a negative effect on earnings, but the effect is statistically insignificant in the specifications that control for actual major. While the model predicts that differences in choice of major between switchers and non-switchers explains only *part* of the earnings gap, the data suggest that differences in choice of major account for the *entire* gap. This casts doubt on the model's assumption that majors generate occupation specific human capital which is rewarded in the labor market.

In addition to addressing the model's prediction about labor market outcomes, the pattern of results in table 6 also implies that high-earnings majors lose students on net. This must be true since the average major switcher earns less than the average non-switcher, but no less than a non-switcher in the same career. This implies that low earnings field gain students on net because major switchers disproportionately switch to them.

VI. Conclusion

Research on the choice of college major has uncovered a number of robust facts about the relationship between major choice, earnings and ability, and has found some consistent patterns in the behavior of major switching. I develop a model of college major choice that incorporates occupation-specific human capital and learning. The model predicts a pattern of major choice in college that is consistent with findings from previous research. I also test the model's predictions using HSB data and the results of these tests generally support the model. I find that students who have higher scores on a high-school test of cognitive ability plan to major in fields with higher average earnings; that receiving low grades in college increases the likelihood that a student will switch majors; and that students who switch majors earn less than students who do not, although this is explained by their choice of majors.

This chapter contributes to the literature on major choice by considering how a simple model of major choice generates predictions consistent with several well known findings. That my results generally support the model suggests that occupation-specific human capital and learning influence students' choices of college major. Some results, however, do not fit the model's predictions. Future research could consider whether focusing on other aspects of major choice (e.g. compensating differentials) might explain these inconsistencies.

Data Appendix

A. Sample Construction

The main sample consists of the subset of High School and Beyond participants who belong to the sophomore cohort, earn a bachelor's degree by the end of the sample period, and have complete information on actual and planned major as well as on control variables. 2,384 students meet these criteria.

For the estimation of earnings equations (reported in table 6) I further restrict the sample to students who make a long-term transition to the labor force during the sample period. To make a long-term transition to the labor force, a student must spend at least three consecutive years mainly working after spending at least 1 year mainly not working. I classify a student as spending the year mainly working if they work full time for a least six months of the year. To work full time, the student's labor market status must be "employed full time" and the student's educational enrollment status must not be "enrolled full time." I also exclude students who are missing any control variables from the earnings equation sample, which consists of 1,775 students.

B. Correspondence between Planned and Actual Field

HSB reports actual major in much greater detail than planned major. I assign actual major codes to planned major groups as follows (listing planned major followed by corresponding CIP codes for actual major): Agriculture: 10, 20, 30, 31; Architecture: 40; Art: 500, 502, 504, 505; Biological Sciences: 260, 261, 262, 263; Business: 60, 61, 62, 63, 70, 71, 80; Communication: 90, 91, 100; Computer Science: 110, 111, 112; Education: 130, 131, 132, 133, 134, 135; Engineering: 140, 141, 142, 143, 144, 150; English: 230, 231, 232, 501; Ethnic Studies: 50, 51, 53; Foreign Languages: 160, 161; Health Occupations: 170, 171, 172, 173, 174; Health Sciences: 180, 181, 185, 186, 187, 188; Home Economics: 190, 191, 192, 200, 201; Interdisciplinary Studies: 300, 301, 302, 303, 304; Mathematics: 271; Music: 503; Philosophy and Religion: 380, 381, 390; Physical Science: 400, 401, 402, 403; Preprofessional: 182, 183, 220, 221; Psychology: 420; Social Sciences: 430, 440, 441, 450, 451, 452, 453, 454, 455, 456; Vocational or Technical: 120, 121, 460, 470, 471, 472, 480, 481, 490, 491; Other: 240, 250, 280, 310, 320, 999.

C. College Placement Council Data

The measure of expected earnings used in table 4 comes from the July, 1982, CPC Salary Survey. The College Placement Council collects data from 161 colleges and universities on job offers made to new college graduates. The July, 1982, averages are based on 51,290 salary offers received by bachelor's degree candidates. I group the starting salary data into 11 major categories, some of which include more than one "actual field." The data do not cover some actual fields, such as education.

References

Berger, Mark C. (1988). "Predicted Future Earnings and Choice of College Major," *Industrial and Labor Relations Review*, 41 (3): 418-429.

College Placement Council. (1983). *CPC Salary Survey: A Study of 1982-1983 Beginning Offers*. Formal Report no. 3. Behlehem, Pa.: College Placement Council.

Davis, James A. (1965). *Undergraduate Career Decisions: Correlates of Occupational Choice*. Chicago: Aldine Publishing Company.

Farber, Henry S. and Robert Gibbons. (1996). "Learning and Wage Dynamics," *Quarterly Journal of Economics*, 111 (4): 1007-1047.

Finnie, Ross and Marc Frenette. (2003). "Earning Differences by Major Field of Study: Evidence from Three Cohorts of Recent Canadian Graduates," *Economics of Education Review*, 22 (2): 179-192.

Freeman, Richard B. (1971). *The Market for College-Trained Manpower*. Cambridge, Mass.: Harvard University Press.

Freeman, Richard B. (1975). "Legal Cobwebs: A Recursive Model of the Labor Market for New Lawyers," Review of Economics and Statistics, 58 (1): 171-179.

Gibbons, Robert and Michael Waldman. (1999). "A Theory of Wage and Promotion Dynamics Inside Firms," *Quarterly Journal of Economics*, 24 (1): 59-107.

Hamermesh, Daniel S. and Stephen G. Donald. (2006). "The Effect of College Curriculum on Earnings: An Affinity Identifier for Non-Ignorable Non-Response Bias," *Mimeo*.

Hoffman, Dennis L. and Stuart A. Low. (1983). "Rationality and the Decision to Invest in Economics," *Journal of Human Resources*, 18 (4): 455-479.

Paglin, Morton and Anthony M. Rufolo. (1990). "Heterogeneous Human Capital, Occupational Choice, and Male-Female Earnings Differences," *Journal of Labor Economics*, 8 (1): 123-144.

Rumberger, Russell W. and Scott L. Thomas. (1993). "The Economic Returns to College Major, Quality and Performance: A Multilevel Analysis of Recent Graduates," *Economics of Education Review*, 12 (1): 1-19.

Werts, Charles E. (1967). "Career Changes in College," *Sociology of Education*, 40 (1): 90-95.

Zarkin, Gary A. (1985). "Occupational Choice: An Application to the Market for Public School Teachers," *Quarterly Journal of Economics*, 100 (2): 409-446.

TABLE 1
Average Characteristics by Planned and Actual College Major

	Number	Switched major	Proportion female	Proportion black	College GPA	Test score	Annual earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Summary statistics by planned college major							
Planned major:							
Agriculture	38	0.605	0.395	0.000	2.89	57.3	22,041
Architecture	29	0.552	0.310	0.034	2.97	60.2	25,103
Art	57	0.561	0.719	0.070	3.08	55.6	16,877
Biological sciences	66	0.621	0.500	0.030	3.15	62.9	21,049
Business	499	0.333	0.555	0.074	2.89	57.8	22,673
Communications	123	0.455	0.585	0.154	2.99	59.3	20,686
Computer science	188	0.633	0.441	0.138	3.01	59.5	24,367
Education	108	0.519	0.806	0.028	3.03	56.7	18,031
Engineering	293	0.543	0.195	0.102	2.90	61.6	26,131
English	53	0.660	0.717	0.094	3.18	61.8	19,819
Foreign languages	23	0.652	0.870	0.130	3.13	60.7	15,630
Health occupations	66	0.894	0.848	0.106	2.98	57.6	21,285
Health sciences	94	0.596	0.862	0.043	3.03	58.1	23,478
Home economics	16	0.563	1.000	0.125	3.01	54.9	18,265
Mathematics	27	0.852	0.519	0.037	3.06	62.5	22,293
Music	28	0.393	0.429	0.071	2.97	59.2	16,420
Philosophy/religion	19	0.684	0.368	0.053	3.21	62.9	45,901
Physical science	70	0.757	0.357	0.071	3.17	62.6	21,342
Preprofessional	265	0.996	0.502	0.121	2.98	60.6	-
Psychology	73	0.740	0.836	0.082	2.96	59.0	21,543
Social Sciences	121	0.529	0.587	0.107	3.04	59.9	22,339
Vocational	40	0.975	0.425	0.075	3.16	58.9	23,312
Other fields	88	0.966	0.682	0.080	2.99	58.2	20,778
All majors	2,384	0.607	0.539	0.089	2.98	59.4	22,678
Standard deviation (across majors)		0.182	0.214	0.040	0.09	2.3	5,954

Notes: Other fields includes students planning to major in ethnic studies, interdiscliplinary studies, or "other."

TABLE 1 (continued)
Average Characteristics by Planned and Actual College Major

	Number (1)	Switched major (2)	Proportion female (3)	Proportion black (4)	College GPA (5)	Test score (6)	Annual earnings (7)
Panel B: Average characteristics by actual college major							
Actual major:							
Agriculture	31	0.516	0.419	0.000	2.91	58.4	21,605
Architecture	19	0.316	0.368	0.000	3.14	61.9	24,493
Art	47	0.468	0.596	0.043	3.09	57.4	18,019
Biological sciences	110	0.773	0.527	0.064	3.21	62.7	19,306
Business	604	0.449	0.520	0.091	2.91	58.2	22,432
Communications	153	0.562	0.634	0.137	2.89	58.3	20,861
Computer science	115	0.400	0.426	0.130	3.01	60.7	27,374
Education	139	0.626	0.871	0.065	3.03	55.6	16,664
Engineering	194	0.309	0.160	0.072	2.98	61.8	29,087
English	107	0.832	0.673	0.047	3.12	61.2	19,749
Foreign languages	24	0.667	0.792	0.083	3.17	61.4	16,288
Health occupations	31	0.774	0.742	0.000	3.07	57.6	24,269
Health sciences	94	0.596	0.894	0.074	2.98	58.1	26,810
Home economics	29	0.759	0.897	0.069	2.81	57.0	18,453
Mathematics	37	0.892	0.514	0.081	3.11	61.8	24,329
Music	27	0.370	0.593	0.074	3.19	59.6	14,811
Philosophy/religion	28	0.786	0.214	0.107	3.23	61.6	19,880
Physical science	57	0.702	0.263	0.070	3.09	62.6	21,885
Preprofessional	5	0.800	1.000	0.000	3.12	56.7	26,694
Psychology	117	0.838	0.658	0.154	2.97	58.8	17,447
Social Sciences	318	0.821	0.487	0.123	2.92	60.1	22,754
Vocational	14	0.929	0.143	0.071	2.92	56.4	20,145
Other fields	84	0.964	0.571	0.048	3.00	59.6	25,750
All majors	2,384	0.607	0.539	0.089	2.98	59.4	22,678
Standard deviation (across majors)		0.201	0.239	0.043	0.12	2.2	3,917

Notes: Other fields includes students who actually earn degrees in ethnic studies, interdiscliplinary studies, or any field that falls into none of the categories for planned major.

TABLE 2
Net Changes by Major

	Planned Major (1)	Actual Major (2)	Net Change (3)	Percent Change (4)
Social sciences	121	318	197	162.8
Business	499	604	105	21.0
English	53	107	54	101.9
Psychology	73	117	44	60.3
Biological sciences	66	110	44	66.7
Education	108	139	31	28.7
Communications	123	153	30	24.4
Home economics	16	29	13	81.3
Mathematics	27	37	10	37.0
Philosophy/religion	19	28	9	47.4
Foreign languages	23	24	1	4.3
Health studies	94	94	0	0.0
Music	28	27	-1	-3.6
Other field	88	84	-4	-4.5
Agriculture	38	31	-7	-18.4
Art	57	47	-10	-17.5
Architecture	29	19	-10	-34.5
Physical science	70	57	-13	-18.6
Vocational	40	14	-26	-65.0
Health occupations	66	31	-35	-53.0
Computer science	188	115	-73	-38.8
Engineering	293	194	-99	-33.8
Preprofessional	265	5	-260	-98.1

Notes: Other fields includes students planning to major in ethnic studies, interdiscliplinary studies, or "other" as well as students who actually earn degrees in ethnic studies, interdiscliplinary studies, or fields that fall into none of the categories for planned major.

TABLE 3
Average Characteristics of Switchers and Non-Switchers

	All (1)	Switchers (2)	Non-Switchers (3)
Proportion female	0.54	0.57	0.51
	(.50)	(.49)	(.50)
Proportion black	0.09	0.09	0.09
	(.28)	(.28)	(.28)
Proportion hispanic	0.08	0.09	0.08
	(.28)	(.29)	(.27)
Father Has College Degree	0.34	0.36	0.32
	(.47)	(.48)	(.47)
Mother Has College Degre	0.19	0.19	0.20
	(.39)	(.39)	(.40)
High Family Income	0.61	0.60	0.61
	(.49)	(.49)	(.49)
Language Other than Engli:	0.20	0.20	0.19
Spoken at Home	(.40)	(.40)	(.39)
High School in South	0.25	0.26	0.24
	(.43)	(.44)	(.43)
Urban High School	0.17	0.18	0.17
	(.38)	(.38)	(.37)
Private High School	0.35	0.36	0.35
	(.48)	(.48)	(.48)
College GPA	2.98	2.98	2.99
	(.46)	(.46)	(.46)
Test score	59.24	59.04	59.48
	(6.35)	(6.69)	(5.87)
Proportion entering labor force	0.84	0.82	0.88
	(.36)	(.39)	(.33)
Monthly earnings	2253	2232	2277
	(1884)	(2382)	(1018)
Total students	2,119	1,184	935

Notes: The numbers in parentheses are standard deviations.

TABLE 4
Relationship Between Ability and Expected Earnings

	Starting Monthly Salary of Planned Major in July, 1982	
	(1)	(2)
Test Score	6.832	6.206
	(1.077)	(1.063)
Female		-170.680
		(13.339)
Black		74.422
		(26.026)
Hispanic		17.852
		(28.560)
Father Has College Degree		-23.769
		(16.015)
Mother Has College Degree		-24.608
		(18.645)
High Family Income		2.376
		(13.942)
Language Other than English		25.742
Spoken at Home		(19.340)
High School in South		16.221
		(15.370)
Private High School		-18.235
		(13.557)
Urban High School		29.275
		(18.011)
Observations	1642	1642
R-squared	0.0235	0.1266

Notes: The numbers in parentheses are White-Huber standard errors.

TABLE 5
Relationship between College GPA and Major Switching

	Switched major (1)	Switched major (2)	Switched major (3)
College GPA	-0.01977	-0.06808	-0.04482
	(.02485)	(.02437)	(.02641)
Female		0.04471	0.03983
		(.02376)	(.02375)
Black		-0.04082	-0.05476
		(.04043)	(.04065)
Hispanic		0.00875	-0.00081
		(.04624)	(.04631)
Father Has College Degree		0.05429	0.05713
		(.02625)	(.02628)
Mother Has College Degree		-0.05165	-0.04595
		(.03064)	(.03078)
High Family Income		-0.01811	-0.01598
		(.02333)	(.02333)
Language Other than English		-0.00358	-0.00335
Spoken at Home		(.03294)	(.03297)
High School in South		0.03358	0.02867
		(.02501)	(.02508)
Private High School		0.02861	0.03083
		(.02290)	(.02291)
Urban High School		0.00982	0.00921
		(.02932)	(.02914)
Test score			-0.00496
			(.00192)
Planned field		yes	yes
	1,862	1,862	1,862
	0.0003	0.1387	0.1415

Notes: The numbers in parentheses are White-Huber standard errors.

TABLE 6
Relationship between Switching College Majors and Monthly Earnings

	log monthly earnings (1)	log monthly earnings (2)	log monthly earnings (3)	log monthly earnings (4)	log monthly earnings (5)	log monthly earnings (6)
Switched major	-0.07510	-0.07250	-0.02256	-0.06396	-0.05567	-0.02200
	(.02172)	(.02186)	(.02222)	(.02138)	(.02107)	(.02139)
Experience	0.02253	0.02190	0.01946	0.01453	0.02081	0.01297
	(.01133)	(.01127)	(.01101)	(.01120)	(.01092)	(.01074)
Age	-0.05670	-0.04836	-0.05373	-0.04816	-0.04969	-0.04777
	(.02777)	(.02785)	(.02571)	(.02741)	(.02690)	(.02555)
Female	-0.20521	-0.19725	-0.15992	-0.16746	-0.15276	-0.12815
	(.02196)	(.02232)	(.02239)	(.02225)	(.02280)	(.02195)
Black	-0.02018	0.00036	-0.00444	-0.00020	0.01132	0.01043
	(.03836)	(.03921)	(.03677)	(.03824)	(.03816)	(.03583)
Hispanic	0.04686	0.06178	0.04937	0.07433	0.06255	0.07018
	(.03691)	(.03720)	(.03596)	(.03536)	(.03794)	(.03625)
College GPA		0.00139	0.02568	0.00059	0.00467	0.01217
		(.02758)	(.02575)	(.02663)	(.02683)	(.02524)
Test score		0.00491	0.00227	0.00387	0.00257	0.00158
		(.00200)	(.00194)	(.00193)	(.00193)	(.00185)
Year	yes	yes	yes	yes	yes	yes
Actual field			yes			yes
Industry				yes		yes
Occupation					yes	yes
Observations	7,741	7,741	7,741	7,741	7,741	7,741
R-Squared	0.0976	0.1000	0.1608	0.1317	0.1576	0.2089

Notes: The dependent variable is ln(annual earnings/number of months working full time). The numbers in parentheses are White-Huber standard errors that correct for clustering at the worker level.

Long Run Health Impacts of Income Shocks: Wine and Phylloxera in 19th Century France

Abhijit Banerjee, Esther Duflo, Gilles Postel-Vinay and Timothy M. Watts

Abstract

This paper provides estimates of the long-term effects on height and health of a large income shock experienced in early childhood. Phylloxera, an insect that attacks the roots of grape vines, destroyed 40% of French vineyards between 1863 and 1890, causing major income losses among wine growing families. Because the insects spread slowly from the southern coast of France to the rest of the country, Phylloxera affected different regions in different years. We exploit the regional variation in the timing of this shock to identify its effects. We examine the effects on the adult height, health, and life expectancy of children born in the years and regions affected by the Phylloxera. The shock decreased long run height, but it did not affect other dimensions of health, including life expectancy. We find that, at age 20, those born in affected regions were about 1.8 millimeters shorter than others. This estimate implies that children of wine-growing families born when the vines were affected in their regions were 0.6 to 0.9 centimeters shorter than others by age 20. This is a significant effect since average heights grew by only 2 centimeters in the entire 19th century. However, we find no other effect on health, including infant mortality, life expectancy, and morbidity by age 20.

I. Introduction

Poor environmental conditions in-utero and early childhood have been shown to have adverse consequences on later life outcomes, including life expectancy, height, cognitive ability, and productivity (Barker, 1992)[1]. Important influences include the disease environment (Almond, 2006), the public health infrastructure (Almond and Chay, 2005), food availability (see e.g. Almond et al. (2007) and Qian (2006) on the great Chinese famine, Rosebloom et al. (2001) and Ravelli et al. (1998) on the Dutch famine of 1944-1945), and even the availability of certain seasonal nutrients (Doblhammer, 2000).

At the same time, evidence from developing countries suggests that young children's nutritional status, not surprisingly, is affected by family income (see e.g. Jensen, 2000, Duflo, 2003). Taken together, these two facts suggest that, at least in poor countries, economic crises may have important long term impacts on the welfare of the cohorts born during these periods.[2] Yet, except for the few papers on famines mentioned above, there is very little evidence establishing a direct link between economic events at birth and adult outcomes. This is perhaps not surprising, since such an analysis requires good data on adult outcomes, coupled with information on economic conditions faced during early childhood. Longitudinal data is often not available over such long time periods, especially in poor countries, except for dramatic events such as famines. As a result, the few existing studies tend to be limited to cohort analyses: For example, Van Den Berg, Lindeboom and Portrait (2006) show that, among cohorts born between 1812

[1] The idea that in-utero conditions affect long run health is most commonly associated and referred to as the "fetal origin hypothesis", and associated to DJP Barker (see e.g. Barker, 1992, 1994). For additional evidence see, among others, Strauss and Thomas (1995), Case and Paxson (2006), Berhman and Rozenzweig (2005) and references therein.

[2] In rich countries this effect may be compensated by the fact that pollutants diminish during economic slumps (Chay and Greenstone, 2003).

and 1912 in the Netherlands, those born in slumps have lower life expectancy than those born in booms. However, a concern with using just time variation is that it might reflect other time-specific elements, such as the quality of public services, the relative price of different nutrients, or even conditions in adulthood. To fill this gap, this chapter takes advantage of the Phylloxera crisis in 19th century France, which, we will argue, generated a negative large income shock that affected different "departments" (roughly similar in size to a US county) in France in different years, combined with the rich data on height and health that was collected by the French Military Administration.

Phylloxera (an insect that attacks the roots of vines) destroyed a significant portion of French vineyards in the second half of the 19th century. Between 1863, when it first appeared in Southern France, and 1890, when vineyards were replanted with hybrid vines (French stems were grafted onto Phylloxera-resistant American roots), Phylloxera destroyed 40 percent of the French vineyards. Just before the crisis, about one sixth of the French agricultural income came from wine, mainly produced in a number of small highly specialized wine-growing regions. For the inhabitants of these regions, the Phylloxera crisis represented a major income shock. Because the insects spread slowly from the southern coast of France to the rest of the country, Phylloxera affected different regions in different years. We exploit this regional variation in the timing of the shock to identify its effects, using a difference-in-differences strategy.

We examine the effect of this shock on the adult height, health, and life expectancy outcomes of children born in years where the Phylloxera affected their region of birth, controlling for region and year of birth effects. The height and health data comes from the military: Every year, all the conscripts (who were a little over 20 at the time of reporting) were measured and the number of young men falling into a number of height categories was reported at the level of the department. The statistics also reported the number of young men who had to be exempted for health reasons, and specified the grounds of exemption. This data is available for 83 departments which are consistently defined over the period we consider.[3] In addition, we use data on female life expectancy at birth constructed from the censuses and reports of vital statistics.

In many ways, France in the late 19th century was a developing country: In 1876 the female life expectancy was 43 years; infant mortality was 22 percent; and the average male height at the age of 20 was 1.65 meters, approximately the third percentile of the American population today. The Phylloxera crisis therefore gives us the opportunity to study the impact of a large income shock to the family during childhood on long-term health outcomes in context of a developing economy. In addition, the crisis had a number of features which makes it easier to interpret the results we get: First, we will show that it was not accompanied by either important changes in migration patterns or by an increase in infant mortality. We therefore do not need to worry about sample selection in those who we observe in adulthood, unlike what seems to happen, for example, in the case of a famine. Second, it did not result in a change in relative prices, which would have meant that it would spill over into areas that do not grow wine (even the price of wine did not increase very much, for reasons we will discuss below). As a result, we can identify regions that were completely unaffected. Finally, the progression of the epidemic was exogenous, as it was caused by the movement of the insects, which is something no one knew how to stop until the late 1880s.

[3] France lost Moselle, Bas-Rhin, and Haut-Rhin to Prussia in 1870, while retaining the Territoire de Belfort, a small part of Haut-Rhin. Moselle, Rhin-Bas, Haut-Rhin, and Belfort are excluded from the analysis.

We find that the Phylloxera shock had a long-run impact on stature: We estimate that children born during a "Phylloxera year" in wine-producing regions are 1.6 to 1.9 milimeters shorter than others. We estimate that this corresponds to a decline in height of 0.6 to 0.9 centimeters for children born in wine growing families in years where their region was affected, a large effect considering that average height grew only 2 centimeters over the entire century. We do not find that children born just before, or just after, the Phylloxera crisis are affected by it, which support the Barker hypothesis of the importance of *in-utero* conditions for long run physical development. However, we also do not find any long-term effect on other measures of health, including morbidity at the age of 20 (measured by the military as well) or life expectancy of women.

The remainder of this chapter proceeds as follow: In the next section, we describe the historical context and the Phylloxera crisis. In section III, we briefly describe our data sources (a data appendix does so in more detail). In sections IV and V, we present the empirical strategy and the results. Section VI concludes.

II. Wine Production and the Phylloxera Crisis

Wine represented an important share of agricultural production in 19th century France: In 1863, the year Phylloxera first reached France, wine production represented about one sixth of the value of agricultural production in France, which made it the second most important product after wheat (table 1). Wine was produced in 79 "departments" out of 89, although it was more than 15% of agricultural production in only 40 of them (we refer to these 40 departments below as the "wine-growing" departments).

Phylloxera is an insect of the aphid family, which attacks the roots of grape vines, causing dry leaves, a reduced yield of fruit, and the eventual death of the plant.

Indigenous to America, the insects arrived in France in the early 1860s, apparently having traveled in the wood used for packaging (though it is possible that it was actually in a shipment of American vines). By the 1860s, the pest had established itself in two areas of France. In the departments on the southern coast, near the mouth of the Rhône, wine growers first noticed the pest's effects in 1863, and there are many records of it in the years 1866-1867. In 1869 the pest also appeared on the west coast in the Bordeaux region. The maps in figures 1A to 1D show the progression of the invasion starting from these two points: From the south, the insects spread northward up the Rhône and outward along the coast. From the west, the insects moved southeast along the Dordogne and Garonne rivers and north to the Loire valley. By 1878, Phylloxera had invaded all of southern France, and 25 of the departments where wine was an important production. It reached the suburbs of Paris around 1885.

During the first years of the crisis, no one understood why the vines were dying. As the Phylloxera spread and the symptoms became well-known, it became clear that the disease posed a serious threat to wine growers, and two of the southern departments (Bouches-du-Rhône and Vaucluse) formed a commission to investigate the crisis. The commission found Phylloxera insects on the roots of infected vines in 1868 and identified the insects as the cause of the dead vines. After experimenting with various ineffective treatments (such as flooding, or treatment with carbon bisulphide), what was to be the ultimate solution was discovered in the late 1880s. This solution required wine growers to graft European vines onto pest-resistant American roots. In 1888, a mission identified 431 types of American vines and the types of French soil they could grow in, paving the road to recovery starting in the early 1890s. Eventually, approximately four fifths of the vineyards originally planted in European vines were replaced with grafted vines.

Figure 2 shows the time series of wine production in France from 1850 to 1908. The decline in the first few years reflects the mildew crisis, which affected the vines before the Phylloxera. After a rapid recovery between 1855 and 1859, the production grew until 1877, by which time more than half of wine growing departments were touched by the Phylloxera. The production fell until 1890, when the progressive planting of the American vines started the recovery.

Table 2 shows the importance of the Phylloxera crisis on wine production. Using data on wine production reported in Gallet (1957), we construct an indicator for whether the region was affected by the Phylloxera. Gallet (1957) indicates the year where the Phylloxera aphids were first spotted in the regions. In most regions, however, for a few years after that, production continued to increase (or remained stable), until the aphid had spread. Since we want to capture the fall in wine production due to the insect, we define the "pre-phylloxera" year as the year before the aphids were first spotted, and the indicator "attained by phylloxera" is equal to one every year when wine production was less than 80% of its level in the "pre-phylloxera" year. It is then turned back to zero after 1890, since the grafting solution had been found by then.[4]

We then run a regression of area planted in vines, the log of wine production, and yield, for the years 1852-1892 on year dummies, department dummies, and an indicator for whether or not the region was touched by the Phylloxera in that year. Namely, we run the following specification:

$$y_{ij}=a\,P_{ij}+k_i+d_j+u_{ij}$$

[4] We experimented with other ways of defining the Phylloxera attack as well, with results that are qualitatively similar (although the point estimates clearly depend on how strong the fall in production has to be before a department is considered to be "affected").

where y_{ij} is the outcome variable (vine grown areas, wine yield, wine production, as well as wheat production) in department i in year j, P_{ij} is an indicator for Phylloxera, k_i and d_j are department and year fixed effects, and u_{ij} is an error term. The standard errors are corrected for auto-correlation by clustering at the department level.

We also run the same specification after controlling for department specific trends:

$$y_{ij}=a\,P_{ij}+k_i+t_{ij}+d_j+u_{ij},$$

where t_{ij} is a department specific trend. The results are shown in table 2, panel B. The yield and the production declined dramatically: according to the specification that controls for a department specific trend, the production was 37% lower and the yield was 42% lower during Phylloxera years.

Unfortunately, yearly data on overall agricultural production or department "income" are not available except for a few departments (Auffret, Hau and Lévy-Leboyer, 1981) so we cannot provide a quantitative estimate of the fall in department "GDP" due to the Phylloxera. However, there are reasons to think that there was no substitution towards other activities, so that the decline in wine production led to a corresponding decline in income in the affected departments. Table 2 shows that the area planted with vines did not decline during the crisis, both because many parcels of land that had been planted with vines would have been ill-suited to all other crops and also because most growers were expecting a recovery. As a result, the decrease in wine production was not compensated by a corresponding increase in other agricultural production: Columns 3 to 5, in the same table, show the results of regressing the production of wheat and the area cultivated on wheat on the Phylloxera indicator, and shows no increase of wheat production compensating the decline in wine production. In the few departments in which the series on agricultural production have been constructed by Auffret et al. (1981), the fall in overall agricultural production does appear to be commensurate with the fall in wine production.

To make up for the shortage of French wine, both the rules for wine imports into France and the making of "piquette" (press cake —the solids remaining after pressing the grape grain to extract the liquids—then mixed with water and sugar) and raisin wines were relaxed. For example, while only 0.2 million hectoliters of wine were imported in 1860s, 10 million hectoliters were imported in the 1880s (for comparison, the production was 24 million hectoliters in 1879). Imports declined again in the late 1890s (Ordish, 1972). As can be seen in figure 1, this kept the price of wine from increasing at anywhere near the same rate as the decrease in production. Price movements thus did little to mitigate the importance of the output shock. Moreover, given the size of the crisis in the most affected regions, farmers could not systematically rely on credit to weather the crisis. In particular, Postel-Vinay (1989) describes in detail how, in the Languedoc region, the traditional system of credit collapsed during the Phylloxera crisis (since both lenders and borrowers were often hurt by the crisis). All of this suggests that Phylloxera was a large shock to the incomes of people in the wine growing regions, and the possibility for smoothing it away was, at best, limited.

III. Data

In addition from the department level wine production data already described, we use several data sources in this paper (the data sets are described in more detail in the data appendix).

First, we assembled a complete department-level panel data set of height reported by the military. Height is widely recognized to be a good measure of general health, and countless studies have shown that it is correlated with other adult outcomes (see Case and Paxson (2006) and Strauss and Thomas (1995) for references to studies showing this in the context of the developing world and, among many others, Steckel (1995), Steckel and Floud (1997), and Fogel (2004) historical evidence on this issue from countries that are now considered developed). France is a particularly good context to use military height data. Since the Loi Jourdan, in 1798, young men had to report for military service in the year they turned 20, in the department where their father lived (all the young men reporting in one department and one year were called a "classe", or military class). They were measured, and since 1836, the data on height was published yearly in the form of the number of young men who fell into particular height categories (the number of categories varied from year to year). These documents also included, for each department and each year, the number of young men exempted and the grounds for exemption (in particular, if they were exempted for disease, the nature of the disease). Using this data, we estimated both parametrically and non parametrically the average height of the 20-year old in each department in each year, as well as the fraction of youth who were shorter than 1.56 meters, the threshold for exemption from military duty (see data appendix for the estimation methods).[5] We also computed the fraction of each military class exempted for health reasons, and created a consistent classification of the disease justifying the exemptions.

The nation-level aggregate data on French military conscripts has been used

[5] Actually 1.56 meters was the highest threshold for exemption that was used---the threshold varied over time.

previously (see, Aron, Dumont and Le Roy-Ladurie (1972); Van Merten (1990), Weir (1997), and Heyberger (2005)). However, this paper is the first to assemble and exploit the data on mean height and proportion of those stunted for all departments and every year between 1872 and 1912.[6] The details of the data construction are presented in the appendix.

Moreover, we also conducted original archival work in three departments to collect military data at the level of the canton (the smallest administrative unit after the "commune", or village) in three wine growing departments (Bouches du Rhône, Gard and Vaucluse). The precise height data was not stored at this level, but the canton level data tells us three things: The fraction of people not inducted into the military for reason of height (which is available only until the cohort conscripted in 1901, since the military did not reject anyone based on height after that year), the fraction rejected for reasons of "weakness", and the fraction put into an easier service for the same reason.

We use the data for the years 1872-1912 (corresponding to years of birth, 1852-1892), which span the Phylloxera crisis, as well as the period before and the recovery, since this is the period for which the military data is the most representative of the height in the population. Starting in 1872, everyone had to report for military service. Starting in 1886, everyone's height was published, even if they were subsequently exempted from the military service. Therefore, the data is representative of the entire population of conscripts from 1886 on, and from 1872 to 1885, we are missing the height data for those exempted for health reasons (15%) or for other reasons (25%). We discuss in the appendix what we assume about the height of those exempted to construct estimates that

[6] With the exception of Postel-Vinay and Sahn (2006), which was written concurrently with this article and exploits the same data set.

are representative of the entire sample, and using a complete individual level dataset available for a sub-sample, we show that these assumptions appear to be valid. The data should therefore be representative of the entire population of young men who presented themselves to the military service examination in a given region.

This sample may be still endogenously selected if Phylloxera led to changes in the composition of those who reported to the military conscription bureau in their region of birth, either because of death, migration, or avoidance of the military service. According to historians of the French military service (e.g. Woloch (1994)), the principle of universal military conscription was applied very thoroughly in France during this period. A son had to report in the canton where his father lived (even if the son had subsequently migrated) and the father was legally responsible if he did not. Avoidance and migration by the son are therefore not likely to be big issues. As a check, we computed the ratio between the number of youths age 15-19 in each department in each census year, and sum of the sizes of the military cohorts in the four corresponding year years (for example, a youth age 19 (resp, 17) in 1856 was a member in the class of 1857 (resp. 1859): The average is 99% and the standard deviation is low (table 1). Moreover, as we will show in table 8, this ratio does not appear to be affected by the Phylloxera crisis.

The main potential sample selection problems that remain are therefore those of differential migration by the fathers and mortality between birth and age 20. We will show in the robustness section that neither of these seems to have been affected by the Phylloxera infestation.

Finally, we use data on number of births and infant mortality (mortality before age 1), from the vital statistics data for each department and each year and two data series constructed by Bonneuil (1997) using various censuses and records of vital statistics: The first one gives the life expectancy of females born in every department every 5 years, from 1806 to 1901. The second gives the migration rates of females, both for the young (20 to 29 year olds) and for the entire population.

IV. Empirical Strategy

A. Department-Level Regressions

Figures 3 and 4 illustrate the spirit of our identification strategy. Figure 3 shows mean height in each cohort of birth for wine-producing regions (where wine represents at least 15% of the agricultural production) and other regions. Wine-growing regions tend to be richer and, for most of the period, the 20 year old males are taller in those regions than in others. However, as Figure 4 shows very clearly, the *difference* in average height between those born in wine-producing regions and those born in other regions fluctuates. The striking fact in this figure is how closely the general trend in the difference in mean height follows the general trend in wine production. Both grow until the end of the 1860s, decline in the next two decades (when the Phylloxera progressively invades France) and increase again in the 1890s, when wine grafting allows production to again rise.

The basic idea of the identification strategy builds on this observation: It is a simple difference in differences approach where we ask whether children born in wine producing departments in years where the wine production is lower due to the Phylloxera, are shorter at age 20 than their counterparts born before or after, relative to those who are born in other regions in the same year. Likewise, we can ask whether they have worse health (as described by the military, lower life expectancy, lower long term fertility, etc.). The difference in differences estimates are obtained by estimating:

(1) $y_{ij}=a\,P_{ij}+ k_i+d_j+u_{ij}$

where y_{ij} is the outcome variable (for example: height at age 20, life expectancy) in departement i in year j, P_{ij} is an indicator for Phylloxera (constructed as explained in section 2, and set to 0 if the department was not a wine producing department[7]), k_i and d_j are department and year fixed effects, and u_{ij} is an error term (following Bertrand, Duflo, and Mullainathan (2004), the standard errors are corrected for autocorrelation by clustering at the department level).

We also run the same specification after controlling for department-specific trends:

(2) $y_{ij}=a\,P_{ij}+ k_i+t_{ij}+d_j+u_{ij}$

Also run is an alternative specification where we use a continuous measure of how much a department is affected by the crisis, instead of the dummy: The measure is constructed by interacting the dummy indicating that Phylloxera has reached the department with the area of vineyard per capita before the crisis (we use average area of vineyards over the years 1850-1869, taken from Gallet, as our pre-crisis measure). An alternative continuous measure is the fraction of wine in agricultural production before the crisis, multiplied by the fraction of the population deriving a living from agriculture.[8]

Finally, in some specifications, we restrict the sample to those departments where wine represented at least 15% of the agricultural production before the crisis. All these regions were affected by the Phylloxera at one point or another, and in these specifications we therefore only exploit the timing of the crisis, not the comparison between regions that may otherwise be different.

[7] I.e. less than 15% of the agricultural income came from wine before the crisis.

[8] We will argue below that this appear to be a correct approximation of the fraction of people who lived in wine growing families.

All these regressions consider the year of birth as the year of exposure. But one could easily imagine that the Phylloxera affects the long-run height (or other measures) even if children are exposed to it later in their lives. We will thus run specifications where the shock variable is lagged by a number of years (so it captures children and adolescents during the crisis). For a specification check, we will also estimate the effect of being born just *after* the crisis.

B. Canton-Level Regressions: Gard, Vaucluse, Bouches-du-Rhône

With regressions run at the department level, one may worry that any relative decline of height during the Phylloxera crisis is due to some other time varying factor correlated with the infestation. This concern is in part alleviated by controlling for department specific trends, but since we are considering a long time period, it is still conceivable that the trends have changed in ways that are different for regions and cohorts affected by Phylloxera, but for reasons that have nothing to do with the disease. We therefore complement this analysis by an analysis performed at the level of the canton: For three departments in Southern France (Gard, Vaucluse, and Bouches du Rhône) we collected data on military conscripts at the canton level from the archives in each of these departments. Data on wine production is not available yearly at this fine a level, but the area of vineyards was collected for the 1866 agricultural enquiry, and is also available at the canton level in the archives. We will combine this measure of the importance of vine production before the crisis with the indicator for the fact that the department as a whole was hit by the disease as a proxy for the impact of Phylloxera on wine-production.

The specification we use for these departments is as follows:

(3) $y_{ijk}=a\ (P_{ij}*V_k)+\ m_k+d_{ji}+u_{ijk.}$

where V_k is the hectares of vineyards in 1866 in canton k, P_{ij} is, as before, a dummy indicating whether department i is affected by the Phylloxera in year j, d_{ji} is a department time year fixed effect and m_k is a canton fixed effect. This specification only compares cantons within the same department, and it fully controls for department times year effects: It thus asks whether young men born in cantons where wine was more important before the crisis suffered more due to the crisis than those born in cantons where it was less important.

IV. Results

A. Height

Table 3 shows the basic results on the height at age 20, at the level of the department. Depending on the specifications, we find that, at age 20, those born in the Phylloxera years are 1.5 to 1.9 milimeters smaller than others, on average. They are also 0.35 to 0.38 percentage points more likely to be shorter than 1.56 centimeters. In panel C and D, we use the continuous measure of exposure to the Phylloxera (hectares of vineyards per capita), first in the entire sample and, in order to check that the results are not driven solely by the contrast between wine-producing and non wine-producing regions, in the sample of departments where wine is at least 15% of agricultural production. We find that, in wine-producing regions, one more hectare of vine per capita increases the probability that a 20 year old male is shorter by 3% and reduces his height by 1 centimeter. If instead, we use data from all the regions, the effects are respectively 1.85% and 7 milimeters.

The best way to scale the effect of the Phylloxera crisis would be to use the fraction of the population living in wine producing households. There is, unfortunately, no data source on this. To proxy it, we use the product of the share of wine in agricultural income before the crisis (in 1862) and the share of the population living in households whose main occupation is agriculture. This would be an overestimate of the fraction of population living in wine producing households[9] if the output per worker was higher for wine than for other agricultural products. Estimates based on the cross-department variation in output per worker and share of wine in agricultural production suggest, however, that the output per worker in wine and non wine production is similar, so that the coefficient of this variable can then be interpreted as the effect of the crisis on wine growing families. We show this specification in panel E (for all regions) and F (for wine producing regions only). The regressions indicate that a child born of a wine producing family during the Phylloxera crisis was 0.5 to 0.9 centimeters smaller by age 20 than he would otherwise have been. This is not a small effect, since heights in France grew only by 2 centimeters in the entire 19th century: In term of stature, being born during the Phylloxera crisis in a wine producing region was equivalent to losing almost half a century worth of growth.

Table 4 shows the results of the specification at the canton level of the 3 wine-producing departments. The results can be directly compared to the results in panel C and D of table 2, since, as in that table, the explanatory variable is the number of hectares of vines divided by the population. One difference is that the data are only available for men born until 1881, and the other is that the threshold for being rejected is not 1.56 meters, but 1.54 meters. This specification also suggests a significant impact of the Phylloxera on height: We find that the probability of being rejected for military service because of height was 1.3% higher for each additional hectare of wine per capita. This is a somewhat

[9] Which would lead to an underestimate of the effect of the Phylloxera when we use this variable.

smaller number than what we found in the department-level specification, but the two numbers are not statistically different, and are of the same order of magnitude. The fact that this specification, which uses a much finer level of variation, provides results that are consistent with the department-level results, is reassuring.

Table 5 investigates the effects of the Phylloxera on various regions and cohorts. Columns 1 to 4 estimate the effect of the Phylloxera on various cohorts. Column 1 is a specification check: We define a "born after Phylloxera" dummy, and confirm that those born immediately after the Phylloxera epidemics are no shorter than those who were born before. This is another element suggesting that the effect is probably not due to an omitted changing trend. Column 2 examines the effect on those who were young children (1 to 2 year olds), or toddlers (2 to 5 year olds) at the time of the epidemic. Interestingly, we find no long-run effects on them. Finally column 3 looks at the effect on those who were teenagers during the crisis, and finds again no effect on them. It seems that the shock had no long lasting effects if it was experienced later in childhood.

Column 4 presents another specification check: We run a specification similar to equation (2) (with the dummy set to 1 for regions where wine represents at least 15% of agricultural production), but we also include a Phylloxera dummy for regions that produce less than 15% of wine. Vineyards in those regions were also affected by the Phylloxera (so that the "Phylloxera" variable can be defined as for all regions) , but we do not expect this to really affect average height, since wine did not affect most of the people in this region. Indeed, the coefficient of the Phylloxera dummy in regions producing little wine is insignificant (and slightly positive).

Column 5 separates the regions that were affected early and those that were affected late (we code as "early" regions where the Phylloxera was first spotted before 1876, the median year at which it reached the regions). Since the disease was progressing slowly from one region to the next, the regions affected later might have been able to anticipate the crisis, and thus avoid a part of the negative impact. This is not what we find: The effect seems to be just as large in regions affected later on.

B. Other Health Indicators

While height is an important indicator of long run health, it is also important to estimate whether the Phylloxera affected more acute indicators of health, both in the short run (infant mortality) and in the longer run (morbidity, life expectancy).

1. Mortality and Life Expectancy

Our next measure of health is the life expectancy of those born during the Phylloxera periods. To study this question, we take advantage of data constructed by Bonneuil (1997), from censuses and corrected vital statistics. Bonneuil constructed with great care life expectancy at birth for women born every 5 years from 1801 to 1901.[10] In column 1 of table 6, we use this variable as the dependent variable, still using the specification from equation (2). We find no impact on life expectancy.

Another (cruder) measure of mortality is given by our data: Since after 1872, everyone age 20 was called for the military service and since, as we will show below, it does not appear that there was selective migration out of the Phylloxera departments, a proxy of survival by age 20 for males is given by the size of the military cohort (*classe*), divided by the number of births in that cohort (column 3). This measure also shows no effect of Phylloxera. We also construct the ratio of the size of the *classe* and the number

[10] Focusing on women avoids biasing the results due to the massive mortality of men during the first world war.

of children who have survived to age 1 (column 4), and again see no effect. Finally, column 6 presents the results on infant mortality (mortality before 1). There again, we find no impact.

2. Military Health Data

To investigate the size of the impact of the Phylloxera infestation on health, we then exploit the data collected by the military on those exempted for health reasons, and the reasons for which they were exempted. We use the same specification as before, with the total number of people exempted, and then the number of people exempted for various conditions, as dependent variables. We use the specification in equation (2) (with department-specific trends).

Column 1 of Table 7 displays a surprising result: The number of young men exempted for health reasons is actually *smaller* for years and departments affected by the phylloxera outbreak. However, the following columns shed light on this surprising result: The incidence of all the precisely defined illnesses (such as myopia, goiter, epilepsie etc.) are in fact unaffected. The only category of health condition that are affected are that of "faiblesse de constitution" (or weakness) and hernia. Weakness is clearly something that is subject to interpretation and one very plausible explanation for what we find is that, faced with the necessity of drafting a fixed number of people for the contingent, the military authorities were stricter with the application of the "weak" category at times where they were rejecting more people for reasons of height. The negative effect on hernia is more puzzling, but could be due to the fact that these cohorts performed less hard physical labor at young ages.

Overall there seems to be no clear evidence of a negative long-run impact of

Phylloxera on any health condition: While it affects long-run height, it seemed not to cause any other physical problems. The conclusion is similar when using "canton" level data (table 4), in the three wine producing regions. In this data, we have the number of people exempted from military service because of weakness. If we consider the years until the year of birth 1881 (class of 1901), there is no impact of the Phylloxera on exemption for weakness. After 1901, however, the military adopted a policy of not rejecting anyone because of height and when we include the post 1901 period, we do see an increase in the number of those exempted for weakness during the Phylloxera period. This suggests that after 1901 some of those who would have been exempted based on their height were now exempted based on being labeled as weak, another indication of the permeability of those two categories.

C. Robustness to Sample Selection

Table 5 presented a number of specification checks confirming the fact that those conscripted 20 years after the department where they reported was affected by Phylloxera, are shorter than others born in other years. However, as we have pointed out above, while a strength of the military data is that it covers a large sample of the male population, there are some reasons to worry about selection biases. The main source of potential selection bias comes from the composition of the military class (the cohort who was drafted in each department in each year), and whether it is representative of all the children born in the department, or of all the children that would have been born in the department in the absence of the Phylloxera.

Selection could arise at several levels. First, there could be fewer births during the Phylloxera period, either due to fewer conceptions or to more stillbirths. The "marginal" (unborn) children could be different than the others, in particular they could have been weaker, had they lived. We regress the number of live births in a year, as well as the ratio of stillbirths to total births in the Phylloxera years in table 6 (columns 2 and 5). There does not appear to be an effect on the number of births.

Second, there could be more infant mortality in the Phylloxera years, thus selecting the "strongest" children. However, we have seen in table 6 that life expectancy and the ratio of the size of the class over the number of births, or the number of survivors by age 1, were unaffected.

Third, children born in Phylloxera years may have been less likely to report for service where they were born. Since they have to report in the place their parents live, this would be due either to avoidance by the child or to migration of the children's parents: Some people may have left the affected regions during these years. Table 8 shows some evidence showing that neither avoidance, nor migration, is likely to be biasing our results: To get a measure of avoidance, we construct the ratio of the size of five subsequent military cohorts on the size of the census cohorts for people age 15-19 (as discussed above, the mean of this ratio is 99%). As shown in column 3 of table 9, this is no different for children who were affected by the Phylloxera, suggesting they were no more likely to avoid military service. Finally, columns 1 and 2 use data constructed by Bonneuil on the migration of women. We use the fraction of women aged 20-29 and aged 20-60 who migrated as dependent variables. There appears to be no effect of the Phylloxera on migration out of the affected departments. This does not mean that people did not migrate at all: In the three departments for which we have canton-level data, we find that find that the classes born in Phylloxera years are smaller in the *cantons* that relied more intensively on wine production (column 4, table 4). This, combined with the fact that we see no migration out of the departments, suggest that some of the families may have migrated out of the wine producing cantons but remained in the department.[11]

V. Conclusion

The large income shock of the Phylloxera had a long-run impact on adult height, most likely due to nutritional deficits during childhood. We estimate that children born in affected regions during the years of the crisis were 0.5 to 0.9 centimeters shorter than their unaffected peers. This is a large effect, considering that height increased only by two centimeters over the period. Similar results are obtained when comparing departments with each other or when using data at a lower level of aggregation, the canton, which reinforces our confidence in their robustness. The effects are concentrated on those born during the crisis, and imply that they suffered substantial nutritional deprivation in-utero and shortly after birth, and those had long term impact on the stature they could achieve.

However, this crisis did not seem to result in a corresponding decline in other dimensions of health, including mortality, even infant mortality. This suggests, as suggested by Deaton, Cutler and Lleras-Muney (2006) that despite the shock to income and corresponding decline in nutrition, health status may have been protected by other factors, for example public health infrastructure (see Goubert (1989) on the importance of clean water in the period).

[11] A regression of the urbanization rate on the Phylloxera dummy (controlling as usual for department dummies, year dummies, and specific department trends) suggests indeed that urbanization progressed significantly faster in Phylloxera years in the affected departments.

Appendix: Data Sources and Construction

This chapter relies on a number of data sources, often assembled in electronic form for the first time. We are pleased to make all the data available for distribution. This appendix details the sources and the procedures followed to construct the data.

1. Wine and Wheat Production

    a. Aggregated Data on Wine Production

The aggregate data on wine and production comes from Statistique Générale de la France, Annuaire Statistique, 1932, Paris, Imprimerie Nationale, 1933, p, 55*-56*.

    b. Department-Level Wine Production and Phylloxera Crisis

The source is: Pierre Galet's *Cépages et Vignobles de France* (Galet, 1957) which gives yearly the area planted in vines and the volume of wine produced in almost every wine-growing department from 1860-1905. We used the yearly series on wine production as is. We computed a measure of pre-phylloxera area planted in wine by computing the average surface planted in vine during the years 1850-1869. We also obtain from Galet the year of the first time the Phylloxera was spotted in every important wine-growing department.

    c. Department-Level Yearly Series on Wheat Production

From 1852 to 1876, the data comes from : Ministère de l'Agriculture et du commerce *Récoltes des céréales et des pommes de terre de 1815 à 1876*, Paris, Imprimerie Nationale, 1878. From 1877 onwards, it comes from Statistique Générale de la France, *Annuaire Statistique*, 1932, Paris, Imprimerie Nationale, 1878--.

    d. Canton-Level Data on Wine Production Pre-Phylloxera

We use the data on wine surface from the Agricultural Inquiry of 1866 (Ministère de l'Agriculture, du Commerce et des Travaux Publics, Enquête agricole, 2° série, enquêtes départementales, 22° et 23° circonscription, Paris, 1867), which are available at the level of the canton.

2. Population data

    a. Department-Level Population

The population (and the agricultural population) of each department was obtained from the 1876 census. Statistique de la France: Résultats généraux du dénombrement de 1876. Paris, Imprimerie nationale, 1878.

b. Canton-Level Population

The data on the population of cantons was taken from the military data (see below). The wine superficie in 1866 was normalized by the canton population in 1872.

c. Births and Deaths

Total number of births, stillbirths, and total number of deaths before age 1 are obtained from the vital statistics, reported annually in "Statistique Générale de la France: Annuaire statistique de la France, Paris, Imprimerie nationale" (various years). We computed infant mortality as the number of deaths before the age of 1 divided by the number of births during the previous year. As Bonneuil (1992) describes in detail, both the data on births and deaths are noisy, and the data has very large outliers. We exclude from the raw series the observations where the infant mortality is in the bottom percentile or the top percentile. This does not remove any department but removes a few very low or very high observations in some departments.

3. Military Height and Health Data

a. Sources: Department-Level Data

The department level data is reported yearly (since 1836) in the yearly publications "Compte rendu sur le recrutement de l'armée". Young men born in a given year, for example 1852, formed the "classe of 1872". They were normally examined early the following year, at the age of between 20 years plus 8 months on average (plus or minus 6 months depending on the date of birth). They were measured and received a medical exam.

Before 1871, not everyone was called for the military service: Each young man received a random number, and people were examined in random order until the size of the "contingent" (active army) needed for the year was filled. It was also possible to exchange with someone else. The size of those exempted for health reasons was not reported.

The data becomes clearer after a 1872 law (following the defeat of Germany) which changed the nature of the military service: Everyone was now called, and replacements were not allowed. The statistics also reported the height of a much larger fraction of the class. There are two periods in the data:

1) From 1886 to 1912: Things are very simple. Every young man was examined, and a summary table reported the distribution of height for the *entire cohort* (as the number of people falling into various bins—the number of bins varied from 10 to 26.
2) From 1872 to 1885:
    - Every young man was called to be examined. From 1872 to 1885, a "classe" was about 300,000 young men. Only about 10,000 per year did not present themselves (3%), and were therefore not measured.
    - Among those who were measured

i. A fraction (about 50% of the cohort) was put in the "contingent", or active service. Their height is recorded in bins in the table "contingent".
ii. A fraction (10%) was put in a category "auxiliaire" (they were asked to do an easier military service). Their height is recorded in bins in the table "auxiliaires", with the same category as the table for the contingent.
iii. A fraction (7%) was put in a category "ajournés" (they were asked to sit out and to go be examined again the following year—fortunately, their data was not aggregated with that of the rest of the contingent the following year). A number of them were "ajournés pour défaut de taille" (too short), and we therefore know that they are smaller than the height threshold in those years (1.54m). A number were "ajournés pour faiblesse de constitution" (weakness) and following Weir (1997) we assume that they are taller than the minimum threshold (otherwise, they would have been "ajournés pour défaut de taille", and that they are otherwise distributed like the rest of the population above this threshold. They represents only a fairly small fraction of the population (about 5%).
iv. A fraction (another 10% of the population) was put in a category "exemptés" (exempted) for health reasons. Still following Weir (1997) we assume that their height has the same distribution as the rest of the population.
v. A fraction (about 25% of the population) was "dispensés", and did not need to do the military service, for reasons other than health (priests, sons of widows, etc. where put in this category). Their height was not reported in the summary descriptive either, and we assume that their height as the same distribution as the rest of the population.

Fortunately, since everyone was measured, it is possible to check some of these assumptions (which this paper is the first to do). Farcy and Faure (2003) collected individual level military data on about 50,000 young men born in 1860 (classe of 1880) in Paris and a few other departments. Figure A1 compares the distribution of height of the those who serve in the active service to that of the "dispensés" (25% of people exempted from the military service for reasons other than health).[12] The distributions are right on top of each other. Figure A2 compares those exempted for health reasons to the entire observed population (dispensés+auxiliaires+ajournés). While the correspondence is less perfect (the distribution of height of the exempted has less mass at the mode), the median of the two distribution are the same (165 centimeters) and the means are also extremely close (165.09 centimeters for the exempted, and 165.26 centimeters for the rest of the population). This suggests that the approximations that are done in this paper are acceptable.

[12] For both figures we estimate the distribution with a kernel density estimator, using an Epanechnikov Kernel and a bandwidth of one centimeter.

- There was a last complication: From 1872 to 1885, the aggregate data was not reported at the level of the department, but at a smaller level (military district). We entered the data at that level and re-aggregated at the department level.

b. Sources: Canton-Level Data

The height data were published after aggregation at the level of the department, but some data was also tabulated at the level of the "cantons". The yearly reports ("Comptes statistiques et sommaires) are kept in the departmental archives, where we collected them for three departments: Gard, Vaucluse, and Bouche du Rhône. The height distribution was not tabulated at that level, but the population of the canton, the number of youth placed in auxiliary service, the number of youth asked to sit out because they were too short, and the number of youth placed in the active service was recorded. Note that we lose this information after the class of 1881, since height stopped being a reason for exemption at this time.

c. Computations Methods

The records on height provide data in a summary form, giving the number of conscripts in height bins of various widths—for example the number of those between 1.54 and 1.63 meters or those between 1.67 and 1.70 meters. Men shorter than the minimum height requirement were exempt from service, and listed as "*défaut de taille*" (lack of height) in the exemption statistics, resulting in left-censored data.[13] The data are in fact double-censored, since the final category includes conscripts above a certain height, 1.73 meters for example. Between 1850 and 1912 the number of categories changes several times, reaching a maximum of 28 from 1903-1912 and a minimum of 9 from 1872 to 1900.

We estimate the departmental distributions of height two ways. First, we assume that the distribution is normal and estimate the parameters of the distribution by maximum likelihood estimation. The log likelihood function is maximized using a simplex search algorithm using starting guesses of 1.66 meters for the mean height and 0.04 meters for standard deviation.[14] Provided the assumption of normality holds, this technique yields efficient estimates of the means and variances of the distributions. We use the mean of the distribution estimated following this procedure.

While height typically follows a normal distribution among adult populations (and the graphs we just discussed indeed appear to be normal), the distribution of height among

[13] The minimum height is 1.56m from 1832 to 1871 and 1.54m from 1872 to 1900. From 1901 on there is no minimum height requirement.

[14] The general form of the log likelihood function for this data is

$$L = c_1 * \ln\left(\Phi\left(\frac{x_1-\mu}{\sigma}\right)\right) + \sum_{i=2}^{n-1}\left(c_i * \ln\left(\Phi\left(\frac{x_i-\mu}{\sigma}\right) - \Phi\left(\frac{x_{i-1}-\mu}{\sigma}\right)\right)\right) + c_n * \ln\left(1-\Phi\left(\frac{x_{n-1}-\mu}{\sigma}\right)\right)$$

where n is the number of height categories, c is the number of observations in each category, and x is a boundary between categories and phi is the standard normal cdf.

French conscripts may differ from a normal distribution. In particular, if the Phylloxera crisis did affect the long term health, it could result in a distorted distribution. In a population with varying levels of initial health, an income shock might therefore not simply shift the distribution; it could have bigger effects at the tails. Imposing a normal distribution on affected regions may obscure these effects.

We therefore also estimate the height distributions non-parametrically. For each department-year cell, the number of observations equals the number of height categories. Each category is assigned the value of the mean of its limits and weighted by its count of conscripts. The extreme lower and upper categories are assigned values one centimeter below and above the censoring points. Using a kernel density estimator, we estimate the PDF evaluated at fifty evenly spaced heights.[15] Riemann summation transforms the points of the PDF into a CDF, and we interpolate to find the deciles of the distribution. We use this estimated data to compute a uniform series for the percentage of individuals measured who are shorter than 1.56 meters (very short), which we use as an alternative dependent variable.

4. Life Expectancy and Migration

The life expectancy and migration data were constructed by Bonneuil (1997) from vital statistics and census data that he very carefully corrected to provide consistent data.

[15] We use a Gaussian kernel with optimal bandwidth approximated by $h = \frac{0.9m}{n^{1/5}}$, where m is the sample variance of the midpoints of the height categories and n is the number of categories.

References

Official Publications

Ministère de l'Agriculture, du Commerce et des Travaux Publics, *Enquête agricole, 2e série, Enquêtes départementales*, Paris, 1867 ; 22e circonscription (Gard, Bouches-du-Rhône) and 23e circonscription(Vaucluse).

Ministère de l'Agriculture et du commerce, *Récoltes des céréales et des pommes de terre de 1815 à 1876*, Paris, Imprimerie Nationale, 1878.

Ministère de la guerre Compte rendu sur le recrutement de l'armée, yearly 1872-1912

Statistique de la France : 2e série, t. III IV X XI and then yearly *Mouvement de la Population*. (Paris, 1853--).

Statistique de la France : 2e série, t. VII and VIII. *Statistique agricole. Enquête agricole de 1852*, 2 vol. (Paris, 1858-60).

Statistique de la France : 2e série, t. XIX. *Résultats généraux de l'enquête décennale de 1862* (Strasbourg, 1868).

Statistique de la France: Resultats généraux du dénombrement de 1876. Paris, imprimerie nationale, 1878

Statistique Agricole de la France, (Algérie et colonies) publiée par le ministre de l'agriculture. *Résultats généraux de l'enquête décennale de 1882* (Nancy, 1887).

Statistique Générale de la France, *Annuaire Statistique*, Paris, Imprimerie Nationale, 1878--.

Statistique Générale de la France, *Annuaire Statistique*, 1932, Paris, Imprimerie Nationale, 1933.

Works Cited

Almond, Douglas. (2006). "Is the 1918 Influenza Pandemic Over? Long-Term Effects of *In Utero* Influenza Exposure in the Post-1940 U.S. Population," *Journal of Political Economy*, 114 (4): 612-712.

Almond, Douglas and Kenneth Chay. (2005). "The Long-Run and Intergenerational Impact of Poor Infant Health: Evidence from Cohorts Born During the Civil Rights Era," Columbia University, Working Paper.

Almond, Douglas, Lena Edlund, Hongbin Li and Junsen Zhang. (2007). "Long-Term Effects of the 1959-1961 China Famine: Mainland China and Hong Kong," Columbia University, Working Paper.

Aron, J. P., P. Dumont and E. Le Roy Ladurie. (1972). *Anthropologie du conscrit français d'après les comptes numériques et sommaires du recrutement de l'armée 1819-1826.* La Haye: Mouton.

Barker, David J. P., ed. (1992). *Fetal and Infant Origins of Adult Disease: Papers.* London: British Medical Journal.

Barker, David J. P. (1994). *Mothers, Babies and Disease in Later Life.* London: British Medical Journal.

Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan. (2004). "How Much Can We Trust Differences in Differences Estimates," *Quarterly Journal of Economics,* 119 (1): 249-275.

Berhman, Jere and Mark Rosenzweig. (2004). "The Returns to Birth Weight," *Review of Economics and Statistics,* 86 (2): 586-601.

Bonneuil, Noël. (1997). *Transformation of the French Demographic Landscape, 1806-1906.* Oxford: Clarendon Press.

Case, Anne and Christina Paxson. (2006). "Stature and Status: Height, Ability, and Labor Market Outcomes," RPDS, Princeton, Working Paper.

Chay, Kenneth and Michael Greenstone. (2003). "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession," *Quarterly Journal of Economics,* 118 (3): 1121-1167.

Cutler, David, Angus Deaton and Adriana Lleras-Muney. (2006). "The Determinants of Mortality," *Journal of Economic Perspectives,* 20 (3): 97-120.

Doblhammer, Gabriele. (2003). "The Late Life Legacy of Very Early Life," Rostock, MPIDR Working Paper WP-2003-030.

Duflo Esther. (2003). "Grandmothers and Granddaughters: Old Age Pension and Intra-Household Allocation in South Africa," *World Bank Economic Review,* 17 (1): 1-25.

Farcy, Jean-Claude and Alain Faure. (2003). *Une Génération de Français à l'Épreuve de la Mobilité. Vers et dans Paris : Recherche sur la Mobilité des Individus à la Fin du XIXe Siècle.* Paris: INED.

Fogel, R. W. (2004). *The Escape from Hunger and Premature Death, 1700-2100.* Cambridge: Cambridge University Press.

Galet, Pierre. (1957-1962). *Cépages et vignobles de France.* Montpellier.

Goubert, Jean-Pierre. (1989). *The Conquest of Water: The Advent of Health in the Industrial Age*. Princeton: Princeton University Press.

Heyberger, L. (2005). *La Rèvolution des Corps*. Strasbourg, Presses Universitaires de Strasbourg.

Jensen, Robert. (2000). "Agricultural Volatility and Investments in Children," *American Econmic Review*, 90 (2): 399-405.

Meng, Xin and Nancy Qian. (2006). "The Long Run Health and Economic Consequences of Famine on Survivors: Evidence from China's Great Famine," Brown University, Working Paper.

Postel-Vinay, Gilles. (1989). "Debts and Agricultural Performance in the Languedocian Vineyard, 1870-1914," in Grantham, G. and Leonard, C. S. eds. *Agrarian Organization in the Century of Industrialization : Europe, Russia, and North America.* Greenwich, Connecticut and London, pages 161-186.

Postel-Vinay, Gilles, and David Sahn. (2006). " Explaining Stunting in Nineteenth Century France," INRA, LEA Working Paper.

Ordish, George. (1972). *The Great Wine Blight*, New York: Scribner.

Ravelli, A. C. J., J. H. P. van der Meulan, R. P. J. Michels, C. Osmond, D. J. P. Barker, C. N. Hales and O. P. Bleker. (1998). "Glucose Tolerance in Adults after Prenatal Exposure to Famine," *Lancet*, 351: 173-177.

Roseboom, T.J., J. H. P. van der Meulen, R. P. J. Michels, C. Osmond, D. J. P. Barker, C. N. Hales and O. P. Bleker. (2001). "Adult Survival after Prenatal Exposure to the Dutch Famine 1944-45," *Pediatric and Perinatal Epidemiology* 15 (3): 220-225.

Steckel, Richard H. (1986). "A Peculiar Population: The Nutrition, Health, and Mortality of American Slaves from Childhood to Maturity," *Journal of Economic History*, 46 (3): 721-741.

Steckel, Richard H. (1995). "Stature and the Standard of Living," *Journal of Economic Literature*, 33(4) 1903-1940.

Steckel, Richard and Roderick Floud. (1997). *Health and Welfare during Industrialization*. NBER, Chicago and London: University of Chicago Press.

Strauss, John and Duncan Thomas. (1995). "Human Resources: Empirical Modeling of Household and Family Decisions," in Berhrman, Jere and T.N. Srinivasan eds. *Handbook of Development Economics*, Amsterdam: North Holland, chapter 34, pages 1885-2023.

Van Meerten, M. A. (1990). "Développement Économique et Stature en France, XIXe-Xxe Siècles," *Annales, Economie, Sociéte, Civilisation*, 45: 755-78.

Van den Berg, Gerard J., Martin Lindeboom and France Portrait. (2006). "Economic Conditions in Life and Individual Mortality," *American Economic Review*, 96 (1): 290-302.

Weir, D. "Economic Welfare and Physical Well-Being in France, 1750-1990," in Steckel, Richard and Roderick Floud eds. *Health and Welfare during Industrialization*, Chicago: University of Chicago Press, pages 161-200.

Woloch, Isser. (1994). *The New Regime; Transformation of the French Civic Order, 1789-1820s*. New York. Norton.

Figure 1A: Phylloxera in 1870

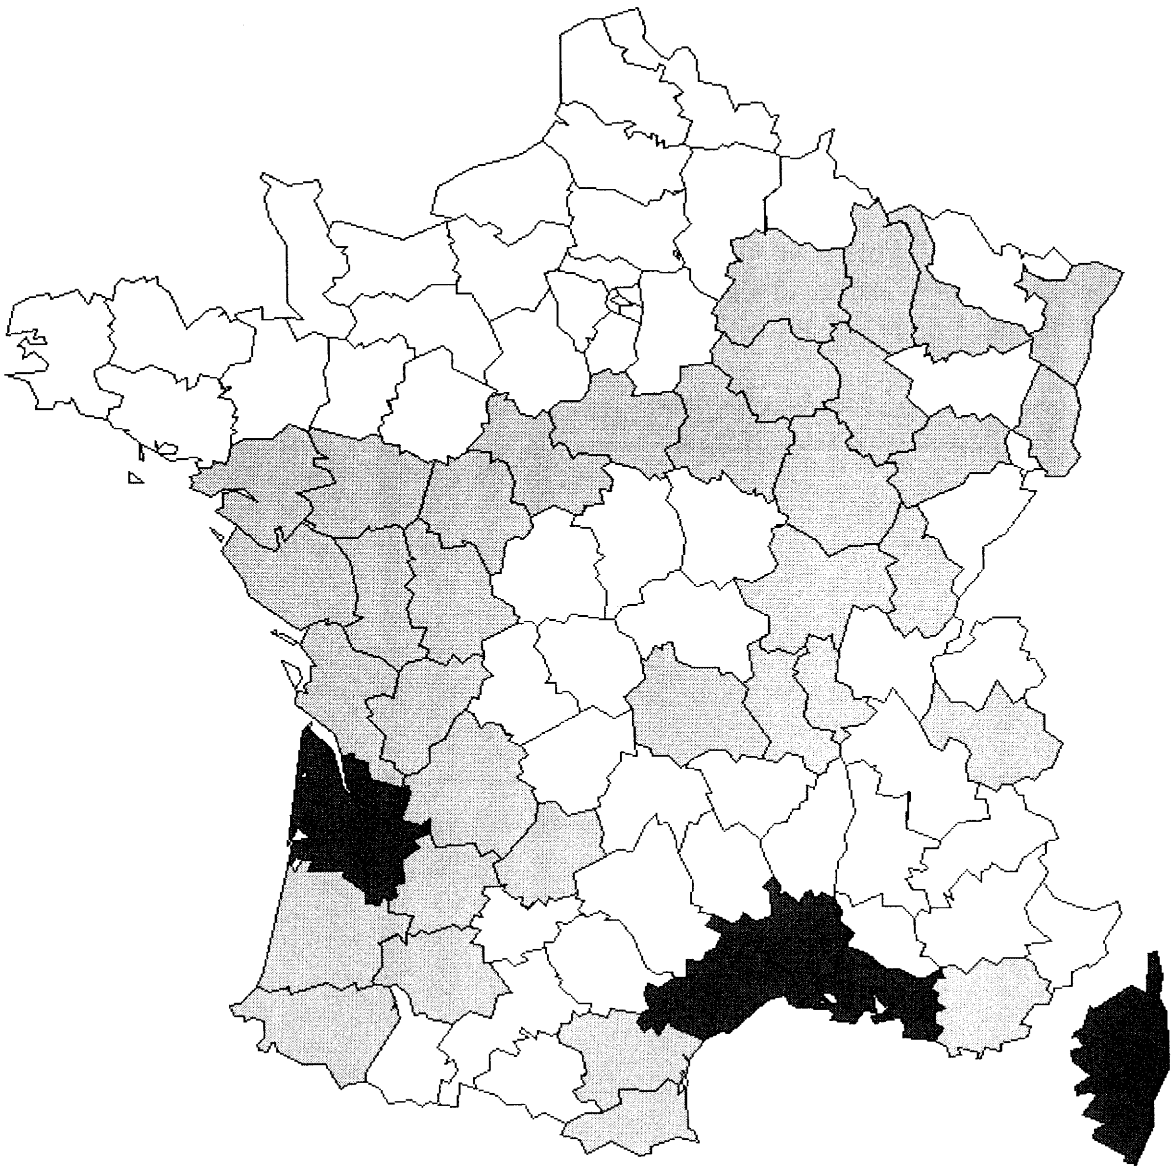

wine-growing departments infected by phylloxera

unaffected wine-growing departments

Figure 1B: Phylloxera in 1875

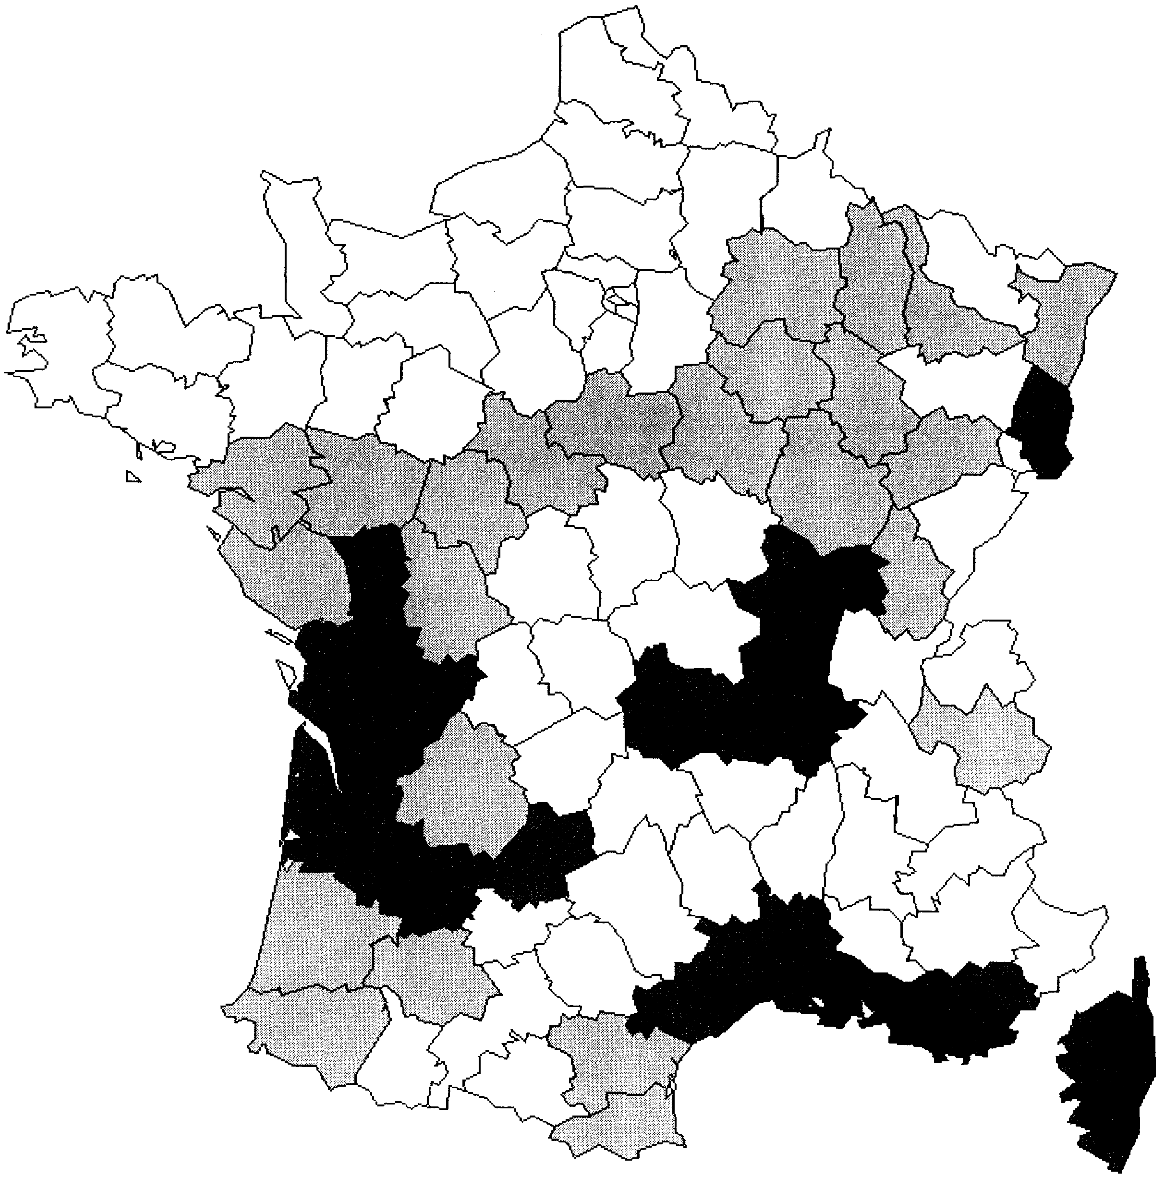

wine-growing departments infected by phylloxera

unaffected wine-growing departments

Figure 1C: Phylloxera in 1880

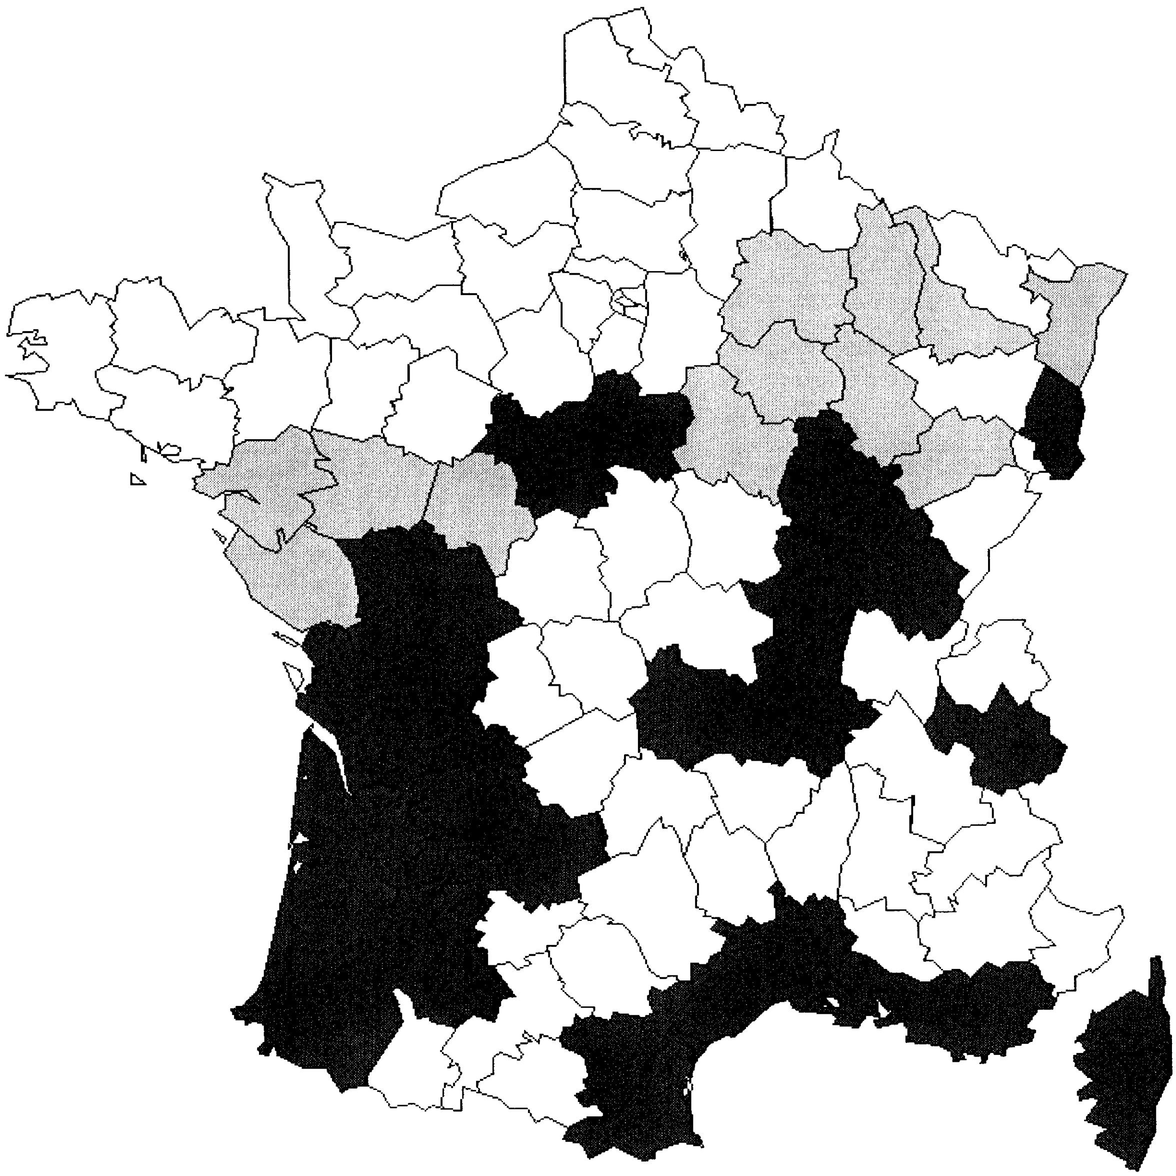

wine-growing departments infected by phylloxera

unaffected wine-growing departments

Figure 1D: Phylloxera in 1890

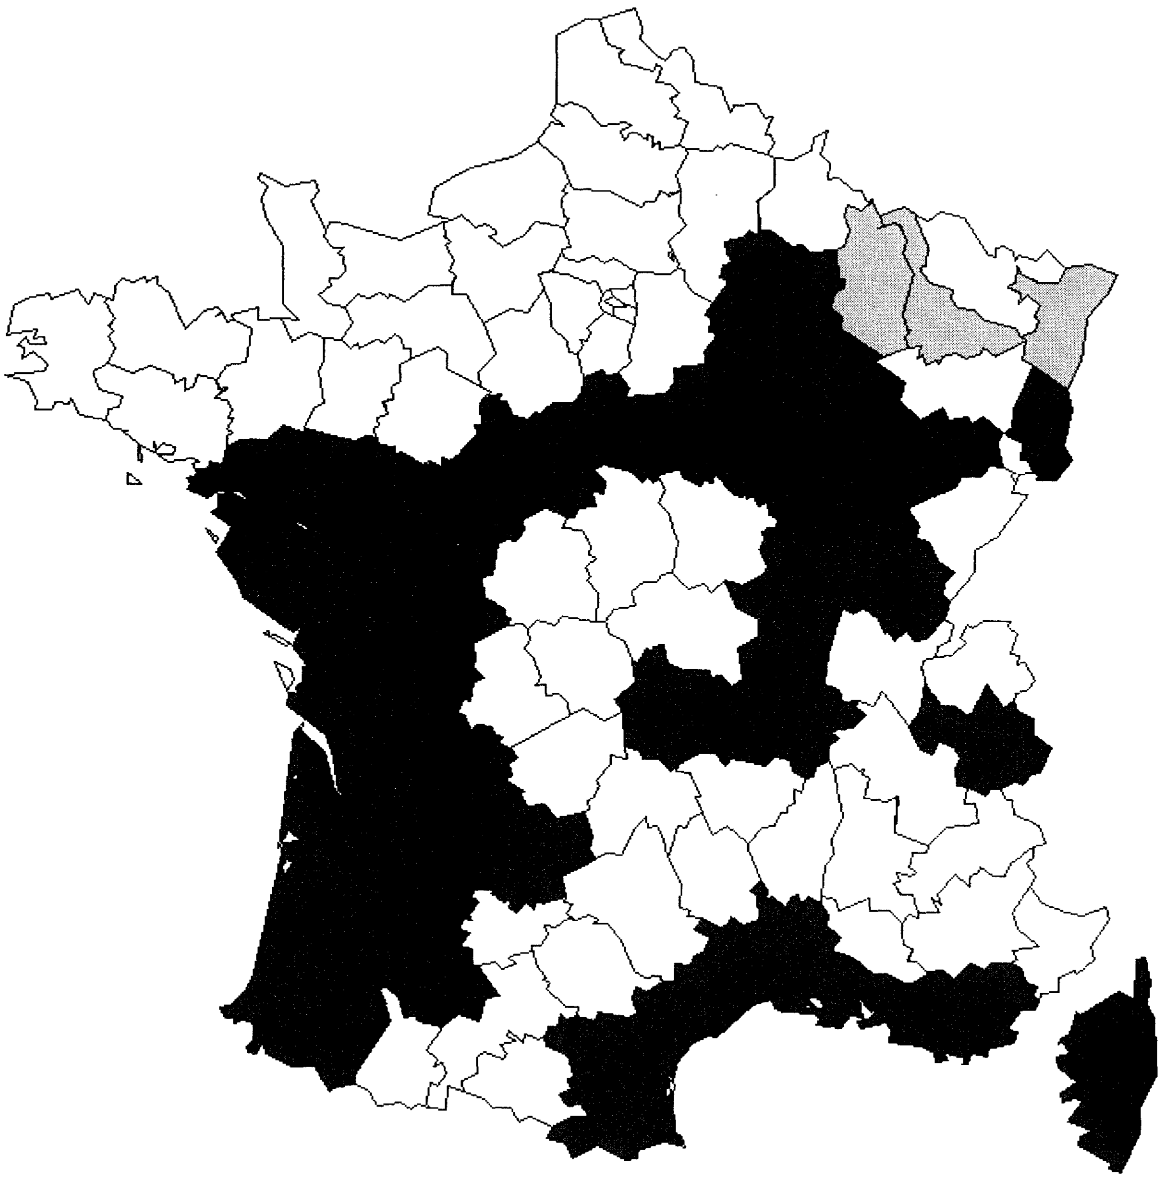

wine-growing departments infected by phylloxera

unaffected wine-growing departments

Figure 2: Wine Production and Wine Price

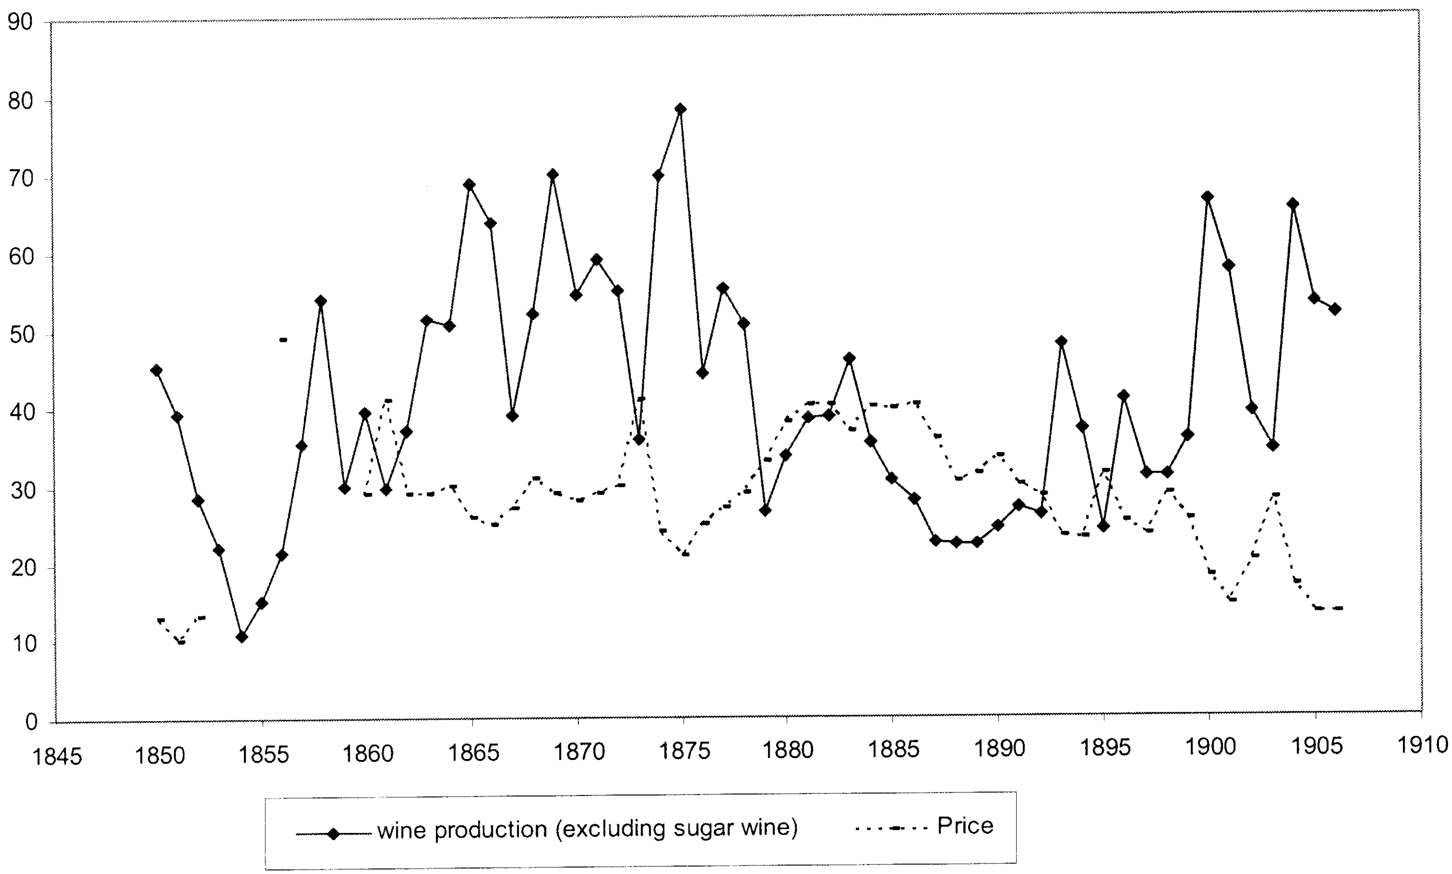

Figure 3: Mean Height Over Time: Wine Producing Regions and Others

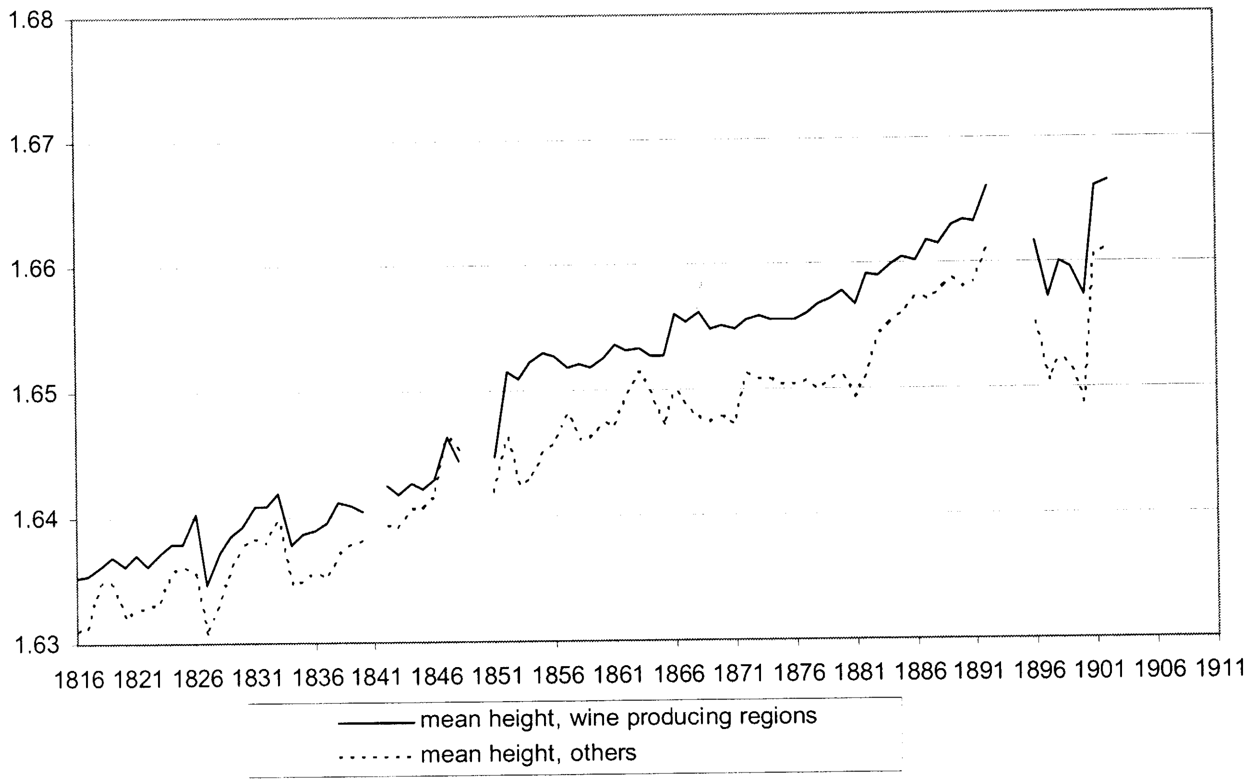

Figure 4: Wine Production and Height Differentials

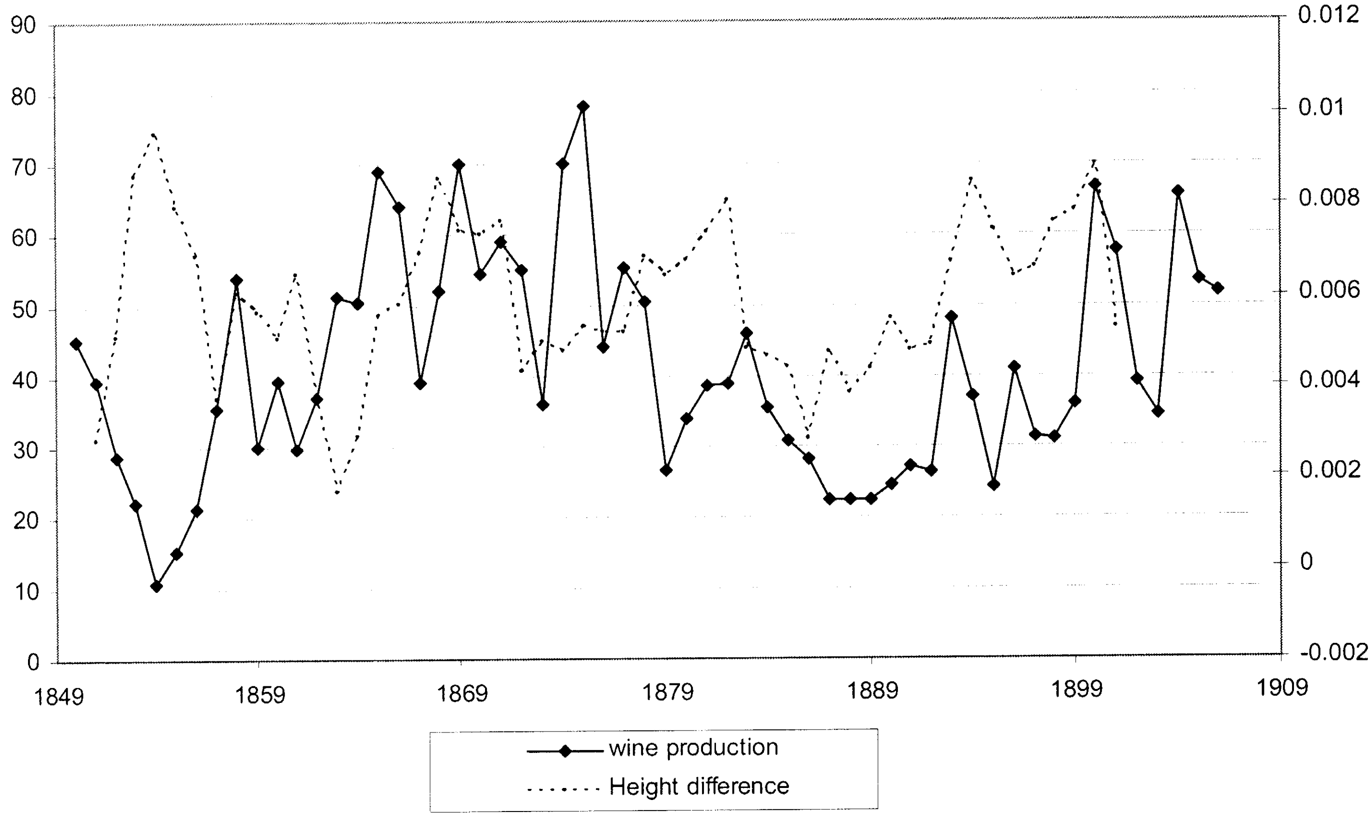

Figure A1: Height Distribution in Individual Data

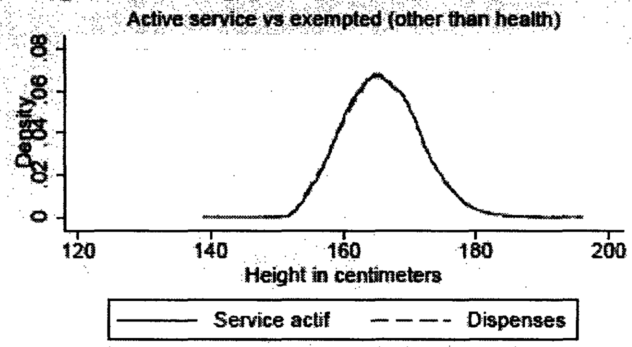

Figure A2: Height Distribution in Individual Data

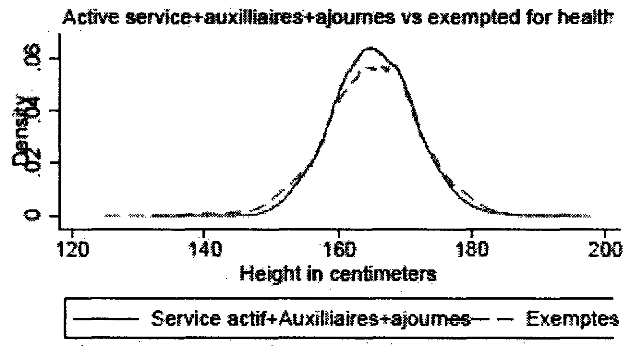

Source: Authors' calculation from data reported in Farcy and Faure (2003)

TABLE 1
Summary Statistics

	Mean	Std. deviation	Observations
	(7)	(8)	(9)
Superficie grown in wine (hectare)	28205	33549	2165
Log wine yield (value per hectare)	2.55	0.83	2153
Log(wine production) (value)	12.03	1.87	3088
Log(wheat production) (value)	13.66	0.89	3209
Share of wine in agricultural production (1863)	0.15	0.14	3649
Share of population working in agriculture	0.58	0.15	3526
Superficie grown in wine per habitant	0.07	0.09	3769
Share of populationin agriculture* share of wine in agricultural Production	0.09	0.08	3526
Number of live births	10751	8593	2795
Share of males surviving untilage 20	0.69	0.23	2683
Share of males surviving until age 20 (conditional on surviving till 1)	0.85	0.29	2599
Share of still births	0.04	0.01	2795
Infant mortality (death before age 1/live birth)	0.17	0.04	2486
Life expectancy (women)	44.34	4.55	630
Net outmigration of youth age 20-29	-0.02	0.05	690
Net outmigration (all)	-0.02	0.02	690
Military class size	3567	2708	3526
Military class/census cohort aged 15-19	0.99	0.27	504
Proportion exempted for health reasons	0.24	0.07	3354
Mean height of military class at 20	1.66	0.01	3526
Fraction of military class shorter than 1.56 meters	0.05	0.02	3526
Height of percentile 10	1.57	0.01	3526
Height of percentile 20	1.59	0.01	3526
Height of percentile 25	1.60	0.02	3526
Height of percentile 50	1.65	0.01	3526
Height of percentile 75	1.69	0.01	3526
Height of percentile 80	1.70	0.01	3526
Height of percentile 90	1.72	0.01	3526

Notes:

1-Except when otherwise indicated, this table presents average by department of the variables used in this paper over the years 1852-1892 (corresponding to the military classes 1872-1912, and years of birth 1852-1892).
2- Data sources are described in details in the data appendix.

TABLE 2
Impact of Phylloxera on Wine Area and Wine Production

	Wine			Wheat		
	log(area)	log(yield)	log(production)	log(area)	log(yield)	log(production)
	(1)	(2)	(3)	(4)	(5)	(6)
A. Controls: Year dummies, departement dummies						
Phylloxera	-0.126	-0.349	-0.399	0.025	-0.009	-0.011
	(.072)	(.055)	(.108)	(.024)	(.026)	(.035)
Observations	2165	2153	3088	1020	1020	3172
B.Controls: Year dummies, departement dummies, departement specific trend						
Phylloxera	-0.090	-0.478	-0.558	0.014	-0.024	0.037
	(.062)	(.071)	(.104)	(.018)	(.025)	(.027)
Observations	2165	2153	3088	1020	1020	3172
C. Sample restricted to wine producing regions (with department specific trend)						
Phylloxera	-0.071	-0.523	-0.570	0.006	-0.010	0.040
	(.062)	(.082)	(.116)	(.016)	(.027)	(.03)
	1058	1051	1519	456	456	1433

Notes:
1-Each column and each panel present a separate regression.
2-All regressions include department dummies and year dummies.
3-Standard errors are corrected for clustering and auto-correlation by clustering at the department level (in parentheses below the coefficient).
4- Importance of wine in production in a department is defined as
share of wine in agricultural production*share of population in an agricultural family.
5-There are fewer data points in this regressions than in the next tables, because the data on wine and wheat production is not available for every department in every year.

TABLE 3

Impact of Phylloxera on Height at Age 20

	Dependent variables	
	Mean height	Fraction shorted than 1.56 meter
	(1)	(2)
A. Year dummies, departement dummies		
Born in phylloxera year	-0.00150	0.00358
	(.00093)	(.00204)
Observations	3485	3485
Department trend	No	No
B.Year dummies, departement dummies, departement trend		
Born in phylloxera year	-0.00188	0.00381
	(.00095)	(.00173)
Observations	3485	3485
Department trend	Yes	Yes
C. Hectare of vine per habitant		
born in phylloxa year *hectare	-0.00753	0.01928
vine per habitant	(.00389)	(.01142)
Observations	3485	3485
Department trend	Yes	Yes
D. Hectare of vines per habitant (wine producing region only)		
born in phylloxera year*hectare	-0.01101	0.03074
vine per habitant	(.00422)	(.01352)
Observations	1558	1558
Department trend	Yes	Yes
E. Share of population in wine growing families		
born in ohylloxera*	-0.00551	0.01198
Importance of wine	(.00392)	(.00985)
Observations	3485	3485
Department trend	Yes	Yes
F. Share of population in wine growing families (wine producing regions)		
Phylloxera*	-0.00901	0.02257
Importance of wine	(.00459)	(.01228)
Observations	1558	1558
Department trend	Yes	Yes

Note:

1-Each column and each panel present a separate regression.

2- The dependent variables are mean height (or proportion shorter than 1.56 meters among a military class in a department and year).

2-"Born in Phylloxera year" is a dummy equal to 1 if the department was affected by phylloxera in the year of birth of a military cohort (see text for the construction of this variable).

3-All regressions include department dummies and year dummies.

4-Standard errors corrected for clustering and auto-correlation by clustering at the department level (in parentheses below the coefficient).

5- The share or population in wine growing families is estimated, as described in the text.

TABLE 4
Effect of Phylloxera on Height, Weakness and Exclusion for Other Reasons
Canton-Level Regression for Three Departments

	Dependent variable: Canton-level means			
	Fraction rejected for size	Fraction rejected for weakness		Class size
	(1)	(2)	(3)	(4)
Hectare of vine per habitant in canton *phylloxera was present in department in year of birth	0.01498	0.00033	0.01664	-64.38
	(.0074)	(.0186)	(.0081)	(37.52)
Observations	2040	2040	2590	2610
Department*year fixed effect	Yes	Yes	Yes	Yes
Canton fixed effects	Yes	Yes	Yes	Yes
Sample	1852-1881	1852-1881	1852-1891	1852-1892
Mean of dependent variable	0.019	0.090	0.10	147
Standard deviation of dependent variable	0.019	0.056	0.06	291

Notes:
1-Data source: archival canton-level data collected by the authors for three departments affected by the phylloxera: Vaucluse, Gard, and Bouches du Rhone.
2-The data set contains canton-level data for these three department for the years 1852-1891.
2-All the regression include separate fixed effects for each year in each department and canton fixed effects.
3- All standard errors (in parentheses below the coefficient) are clustered at canton-level.

TABLE 5
Who is Affected?

	Dependent variables									
	Mean Height at 20					Fraction shorter than 1.56 meters at 20				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Born in phylloxera year	-0.00220	-0.00177	-0.00170	-0.00301	-0.00186	0.00550	0.00305	0.00349	0.00302	0.00221
	(.00119)	(.00093)	(.00091)	(.00121)	(.00084)	(.00239)	(.00181)	(.00173)	(.00229)	(.0017)
Born 1 to 5 years after phylloxera	-0.00108					0.00563				
	(.00171)					(.00626)				
Born 1 to 2 year before phylloxera		0.00000					-0.00183			
		(.00091)					(.00127)			
Born 2 to 5 years before phylloxera		0.00079					-0.00384			
		(.00119)					(.0011)			
Born in phylloxera year, region affected early										0.00705
										(.00218)
Born in phylloxera year, region producing little wine				0.00127					0.00088	
				(.00081)					(.00176)	
Teenager during phyloxera			0.00127					-0.00222		
			(.00083)					(.00104)		
Number of observations	3485	3485	3485	3485	3485	3485	3485	3485	3485	3485

Notes:
1-All regressions include year of birth dummies, department dummies, and department specific trends.
2-All standard errors (in parentheses below the coefficient) are accouting for clustering and autocorrelation by clustering at the department level.

TABLE 6

Impact on Fertility, Mortality and Life Expectancy

	Life expectancy (women)	Live births	Class size /live births in birth year	Class size /survivors at age 1	Still births/all births	Infant mortality (before age 1)
	(1)	(2)	(3)	(4)	(5)	(6)
Born in phylloxera year	0.5087	-39	-0.0025	-0.0082	-0.0004	-0.0033
	(.3604)	(76)	(.0062)	(.0087)	(.0005)	(.0024)
Number of observations	622	2763	2651	2568	2763	2461

Notes:

1-All regressions include year of birth dummies, department dummies, and department specific trends.

2-All standard errors (in parentheses below the coefficient) are accouting for clustering and autocorrelation by clustering at the department level.

3- Data on life expectancy for women used in Column 1 is from Bonneuil (1997).

4-Data on live births, infant mortality, and still births obtained from vital statistics.

5-Data on class size obtained from mility record. A class observed in year t was born in year t-20.

TABLE 7

Effect on Health Outcomes, Military Data

	Exempt due to health	Myopia	Goiter	Hernia	Spinal problem	Epilepsy	Low IQ	Feeble	Blind	Deaf
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Born in	-0.0081	-0.00012	-0.00010	-0.00092	-0.00025	0.00001	0.00020	-0.00218	0.00008	0.00020
Phylloxera year	(.0044)	(.00014)	(.00014)	(.00032)	(.00028)	(.00011)	(.00023)	(.00077)	(.00012)	(.0001)
Number of observations	3315	3485	3485	3485	3485	3485	3485	3485	3485	3485

Notes:

1-All regressions include year of birth dummies, department dummies, and department specific trends.

2-All standard errors (in parentheses below the coefficient) are accouting for clustering and autocorrelation by clustering at the department level.