A CONSTRUCTION OF WIENER MEASURE USING CARATHÉODORY'S THEOREM

BENJAMIN PITTMAN-POLLETTA

Contents

1. Introduction

When a microscopic particle such as a pollen grain is suspended in a solution it moves erratically and rapidly. This motion is due to collisions of the particle with the molecules of the solvent, which move due to their intrinsic kinetic energy (temperature). It is called Brownian motion, after the English botanist Robert Brown. Brownian motion was discovered by Leeuwenhoek in the 15th century, but it was Brown's 1828 research that showed that the phenomenon was not exclusive to organic matter. In 1905, Einstein and Smoluchowski independently published papers showing that the statistical properties of Brownian motion could be accounted for by molecular collisions; Einstein went on to derive the dimensions of water molecules from these statistical properties. This year marks the 100th anniversary of that discovery, which confirmed for many scientists the existence of atoms, and Brownian motion continues to find applications in many fields [1].

Date: May, 2005.

The paths of particles experiencing Brownian motion are exceptionally erratic as well as being random, and their average behavior is not easy to describe. If a pollen grain is placed in the middle of a petri dish of water, how likely is it that the pollen grain will come back to the middle of the petri dish? How likely is it that the pollen grain will reach the edge of the petri dish? If it does, how long will it take? To answer questions like these without performing thousands of difficult experiments on pollen grains, it is useful to develop a mathematical model of Brownian motion that allows such probabilities to be measured. Applied mathematician Norbert Wiener was the first to do this. His mathematical model of Brownian motion is called the abstract Wiener process, and Wiener measure is a probability measure on the space of all Brownian paths. It can be used to calculate the probability that Brownian paths will exhibit certain behaviors.

As one of the few measures with a concrete physical interpretation, Wiener measure is a valuable pedagogical tool, particularly for students in applied mathematics. Unfortunately, the standard construction of Wiener measure involves concepts from functional analysis, including the theory of linear functionals and the Riesz Representation theorem. Most students of analysis don't encounter these ideas until they are already comfortable with measures, and as a result some of the educational value of Wiener measure is lost. This paper is an attempt to construct a version of Wiener measure from more basic principles. We restrict ourselves to the basics of measure theory, and so in constructing the measure we must rely on Carathéodory's theorem. This theorem is the most general and basic way to construct measures, and it is used to construct Lebesgue measure. Our construction of a version of Wiener measure closely parallels the construction of Lebesgue measure.

During our construction of the measure, we found it necessary either to restrict ourselves to functions from $\mathbb{Q} \cap [0, \infty)$ to $\mathbb{R}$, or to restrict our measure to the σ-algebra $\bigotimes_{t \in [0,\infty)} \mathcal{B}_{\mathbb{R}}$, that is, the product over $[0, \infty)$ of the Borel σ-algebra on $\mathbb{R}$. We chose the former in the hopes that the measure we constructed could then be suitably extended to $\mathbb{R}^{[0,\infty)}$. This turns out to be the right choice, because it allows us to prove that our measure is supported on the Hölder continuous functions of exponent $\alpha < \frac{1}{2}$, a result that is true for classical Wiener measure. This, in turn, allows us to extend our measure to the space of functions from $[0, \infty)$ to $\dot{\mathbb{R}}$.

The outline of the paper is as follows. In the second section, preliminaries from topology, measure theory, and probability are presented, with an emphasis on product topologies and product σ-algebras. Carathéodory's theorem is also stated. In the third section, the Wiener process is derived from physical principles. In the fourth section, Wiener measure is constructed using Carathéodory's theorem. In the fifth section, our Wiener measure is shown to be concentrated on Hölder continuous paths. In section six, we discuss how our measure can be extended to the full space of Brownian paths. Some conclusions are offered in the final section. To my knowledge, Lemma 5.4 and most of the ideas in sections four and six are original. The rest of section five follows [3] quite closely; most of the other material follows [2].

2. Preliminaries

What follows relies heavily on topology and measure theory, with a dash of probability. We provide a brief review of these subjects, with an emphasis on the properties of product topologies and product σ-algebras. Proofs will not be given. The reader who is already conversant in graduate-level analysis is encouraged to skip the sections on topology and measure theory; we will refer to them as needed. A number of concepts from probability theory will be used to derive the Wiener process in Section 3, but they will not be used in the rest of the paper.

2.1. **Topology.**

2.1.1. *Basic Concepts.* Almost all concepts in analysis are formulated using limits. The usual $\epsilon - \delta$ definition of a limit applies only in the context of metric spaces, where the distance between any two points can be calculated. In order to extend analysis to more general spaces, a more general concept of how close two points are is necessary. This is the motivation for topology. In a topological space, a family of subsets called "open sets" is specified. Two points are then said to be close together if they inhabit many of the same open sets. The family of open sets is called a topology. In words, a topology is a family of subsets closed under arbitrary unions and finite intersections.

Definition 2.1. Given a set X, a topology τ on X is a family of subsets of X such that

(i) $\emptyset \in \tau$, $X \in \tau$,
(ii) If $\{U_\alpha\}_{\alpha\in A} \subset \tau$, then $\bigcup_{\alpha\in A} U_\alpha \in \tau$,
(iii) If $\{U_i\}_{i=1}^n \subset \tau$, then $\bigcap_{i=1}^n U_i \in \tau$.

The elements of the topology are the open sets, and a space together with a topology defined on it is called a topological space. When referring to a topological space, we will often omit the topology and refer only to the space.

In general, an open set can be very complicated - for example, in the plane, any connected set without a boundary is open, no matter how bizarrely it is shaped. Rather than attempting to describe all the open sets in a topology τ, we can specify a generating family $\mathcal{E}$ of sets from which all elements of τ can be created via unions (arbitrary) and intersections (finite).

Definition 2.2. If $\mathcal{E} \subset \mathcal{P}(X)$ for some space X, then the unique smallest topology $\tau(\mathcal{E})$ containing $\mathcal{E}$ is the topology generated by $\mathcal{E}$, and $\mathcal{E}$ is said to be a subbase for τ.

The set of all open disks is a subbase for the usual topology on $\mathbb{R}^2$, $\tau_{\mathbb{R}^2}$. While a subbase generates only one topology, a given topology may have many subbases; the open cubes, the open triangles, and the open moon-shaped sets all also generate $\tau_{\mathbb{R}^2}$.

2.1.2. *Compactness.* In particular, we will be interested in topologies exhibiting a property called compactness. An open covering of a set $E \subset (X, \tau)$ is a family of sets $\{U_\alpha\}_{\alpha\in A} \subset \tau$ such that $E \subset \bigcup_{\alpha\in A} U_\alpha$. Intuitively, it is a covering of E by open sets.

Definition 2.3. A subset E of a topological space (X, τ) is compact if whenever $\{U_\alpha\}_{\alpha \in A}$ is an open covering of E, there is a finite subset B of A such that

$$E \subset \bigcup_{\alpha \in B} U_\alpha.$$

$\{U_\alpha\}_{\alpha \in B}$ is called a finite subcovering of $\{U_\alpha\}_{\alpha \in A}$, so E is compact if every open covering of E admits a finite subcovering.

A topological space (X, τ) is said to be compact if X itself is compact.

2.1.3. *Product Spaces.* Given a collection of topological spaces, it is possible to create a new topological space from their set-theoretic product. Perhaps the most useful of the topologies that can be put on this product is named the product topology; the resulting space is called the product space.

Definition 2.4. If $\{(X_\alpha, \tau_\alpha)\}_{\alpha \in A}$ is a collection of topological spaces, and $X = \prod_{\alpha \in A} X_\alpha$, then the product topology on X is the topology generated by the sets

$$V = \pi_\alpha^{-1}(U_\alpha)$$

where π_α is the projection from X onto X_α, and $U_\alpha \in \tau_\alpha$.

If, for all α, $X_\alpha = Y$ for some topological space Y, then X above can be viewed as the set of all functions from A to Y. In this case, the product topology is also the topology of pointwise convergence; that is, if $\{f_n\}_1^\infty \subset X$ converges to $f \in X$ in the product topology, $\{f_n\}_1^\infty$ converges to f pointwise as well, and vice-versa.

If each space X_α is a metric space, with associated metric ρ_α, then there is another topology that can be put on the product space, called the topology of uniform convergence. In the case that the product space is a function space, convergence in this topology is equivalent to uniform convergence.

Definition 2.5. If $\{(X_\alpha, \rho_\alpha)\}_{\alpha \in A}$ is a collection of metric spaces, and $X = \prod_{\alpha \in A} X_\alpha$, then the topology of uniform convergence on X is the topology generated by the open balls of the metric

$$\rho(x, y) = \sup_{\alpha \in A} \rho_\alpha(\pi_\alpha(x), \pi_\alpha(y)).$$

Here $x, y \in X$, and π_α is the projection from X onto X_α (so $\pi_\alpha(x)$ is the α^{th} coordinate of x).

In general, it is hard to show that a space is compact. The following theorem says that showing the compactness of product spaces is easy. Unfortunately, it is as difficult to prove as it is useful. The proof requires Zorn's Lemma, and can be found in many other papers and books.

Theorem 2.6 (Tychonoff's Theorem). *If $\{X_\alpha\}_{\alpha \in A}$ is a collection of compact topological spaces, then $X = \prod_{\alpha \in A} X_\alpha$, equipped with the product topology, is compact.*

2.1.4. *Continuous Functions.* We would be remiss if we did not mention continuous functions in this brief review. A continuous function is one that preserves topological properties.

Definition 2.7. Given topological spaces (X, τ_X) and (Y, τ_Y), a function $f : X \to Y$ is continuous if for each $U \in \tau_Y$, $f^{-1}(U) \in \tau_X$.

2.2. **Measure Theory.**

2.2.1. *σ-Algebras.* The geometric concepts of length, area, volume, and their higher-dimensional generalizations (which are all subsumed by the handy term hypervolume) are crucial to many concepts in analysis, including Riemann integration. Unfortunately, they are defined for a very limited class of subsets of $\mathbb{R}^n$. Measure theory extends these concepts to much larger classes of subsets of arbitrary spaces using set functions called measures. The classes of subsets on which measures are defined are called σ-algebras. In words, a σ-algebra is a family of sets that is closed under countable unions and complements.

Definition 2.8. Given a space X, a σ-algebra on X is a nonempty family Σ of subsets of X such that

(i) If $\{E_i\}_{i=1}^\infty \subset \Sigma$, then $\bigcup_{i=1}^\infty E_i \in \Sigma$,
(ii) If $E \in \Sigma$, then $E^c \in \Sigma$.

Note that, thanks to de Morgan's laws, σ-algebras are also closed under countable intersections. Also, if $E \in \Sigma$, $E \cup E^c = X \in \Sigma$ and $X^c = \emptyset \in \Sigma$. A space together with a σ-algebra defined on it is called a measurable space, and the elements of the σ-algebra are called measurable sets. As with topological spaces, we will often neglect to mention the σ-algebra when talking about a measurable space.

An arbitrary measurable set can be even more complicated than an arbitrary open set (see [2] for a good discussion), so once again we rely on specifying a generating family of sets, rather than attempting to characterize a general measurable set.

Definition 2.9. If $\mathcal{E} \subset \mathcal{P}(X)$ for some space X, then the unique smallest σ-algebra $\mathcal{M}(\mathcal{E})$ containing $\mathcal{E}$ is called the σ-algebra generated by $\mathcal{E}$.

When we are in a topological space (X, τ), it is often useful to have a σ-algebra which incorporates the underlying topology. $\mathcal{M}(\tau)$, the σ-algebra generated by τ, is called the Borel σ-algebra on X, and is denoted $\mathcal{B}_X$. It is useful to know that $\mathcal{B}_\mathbb{R}$, the Borel σ-algebra on $\mathbb{R}$, is generated by the open intervals, the closed intervals, and the (right or left) half-open intervals.

The next lemma is handy in proving σ-algebras equivalent.

Lemma 2.10. *If $\mathcal{E} \subset \mathcal{M}(\mathcal{F})$, then $\mathcal{M}(\mathcal{E}) \subset \mathcal{M}(\mathcal{F})$.*

Just as a product of topological spaces can be made into a topological space, a product of measurable spaces can be made into a measurable space, and in the same way.

Definition 2.11. If $\{(X_\alpha, \Sigma_\alpha)\}_{\alpha \in A}$ is a collection of measurable spaces, and $X = \prod_{\alpha \in A} X_\alpha$, then the product σ-algebra on X is the σ-algebra generated by the sets

$$\pi_\alpha^{-1}(E_\alpha)$$

where π_α is the projection from X onto X_α, and $E_\alpha \in \Sigma_\alpha$. We denote this σ-algebra by $\bigotimes_{\alpha \in A} \Sigma_\alpha$.

If we have generating sets for each Σ_α above, we can construct a generating set for the product σ-algebra as well.

Proposition 2.12. *Let $\{(X_\alpha, \Sigma_\alpha)\}_{\alpha\in A}$ be a collection of measurable spaces, and for each $\alpha \in A$, let $\mathcal{E}_\alpha$ be a generating set for Σ_α. Then*

$$\mathcal{E} = \{\pi_\alpha^{-1}(E_\alpha) : E_\alpha \in \mathcal{E}_\alpha, \alpha \in A\}$$

generates $\bigotimes_{\alpha\in A} \Sigma_\alpha$.

In general, the product σ-algebra is contained in the Borel σ-algebra generated by the product topology, but the two are not equivalent. For countable products of certain types of spaces (in particular, second-countable spaces), it is straightforward to show that the product σ-algebra is the σ-algebra generated by the product topology; we will use this fact in Section 4.

2.2.2. *Measures.* We now have all that we need to define measures.

Definition 2.13. Let (X, Σ) be a measurable space. A measure on Σ is a function $\mu : \Sigma \to [0, \infty]$ such that

(i) $\mu(\emptyset) = 0$,
(ii) If $\{E_i\}_{i=1}^\infty$ is a sequence of disjoint sets in Σ, then

$$\mu\left(\bigcup_{i=1}^\infty E_i\right) = \sum_{i=1}^\infty \mu(E_i).$$

A measurable space, when equipped with a measure, becomes a measure space. We note here that a set function on a σ-algebra is called finitely additive if whenever $\{E_i\}_{i=1}^n$ is a sequence of disjoint sets in Σ, $\mu\left(\bigcup_{i=1}^n E_i\right) = \sum_{i=1}^n \mu(E_i)$.

Several properties of measure spaces of which we will make frequent use are collected below.

Theorem 2.14. *Let (X, Σ, μ) be a measure space.*

a. *(Monotonicity) If $E, F \in \Sigma$ and $E \subset F$, then $\mu(E) \le \mu(F)$.*
b. *(Subadditivity) If $\{E_j\}_{j=1}^\infty \subset \Sigma$, then $\mu(\bigcup_{j=1}^\infty E_j) \le \sum_{j=1}^\infty \mu(E_j)$.*
c. *(Continuity from below) If $\{E_j\}_{j=1}^\infty \subset \Sigma$, and $E_j \subset E_{j+1}$ for all j, then $\mu(\bigcup_{j=1}^\infty E_j) = \lim_{n\to\infty} \mu(E_j)$.*
d. *(Continuity from above) If $\{E_j\}_{j=1}^\infty \subset \Sigma$, $\mu(E_1) < \infty$, and $E_j \supset E_{j+1}$ for all j, then $\mu(\bigcap_{j=1}^\infty E_j) = \lim_{n\to\infty} \mu(E_j)$.*

2.2.3. *Carathéodory's Theorem.* The problem that now faces us is that we do not know that any measures exist. More to the point, we do not know whether or not there exist measures with the properties we are interested in. For example, we would like a measure that extends the concept of length to all the measurable sets in $\mathcal{B}_\mathbb{R}$; that is, we would like to find a measure μ on $\mathcal{B}_\mathbb{R}$ such that $\mu((a, b)) = b - a$. Carathéodory's Theorem tells us exactly when and how we can construct measures to our specifications. First we will need a few definitions.

Definition 2.15. Given a space X, an algebra on X is a nonempty family $\mathcal{A}$ of subsets of X such that

(i) If $\{A_i\}_{i=1}^n \subset \mathcal{A}$, then $\bigcup_{i=1}^n A_i \in \mathcal{A}$,
(ii) If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$.

As with σ-algebras, these properties ensure that $\mathcal{A}$ is closed under finite intersections, and that $\emptyset, X \in \mathcal{A}$. The restriction to finite unions and intersections makes it much easier to characterize the elements of an algebra succintly, but it is still sometimes useful to speak of generating families for algebras; the definition is the same as for topologies and σ-algebras.

Definition 2.16. Let $\mathcal{A}$ be an algebra on a set X. A premeasure is a function $\mu_0 : \mathcal{A} \to [0, \infty)$ such that

(i) $\mu_0(\emptyset) = 0$,
(ii) If $\{A_i\}_{i=1}^{\infty}$ is a sequence of disjoint sets in $\mathcal{A}$ and $\bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$ as well, then

$$\mu_0\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu_0\left(A_i\right).$$

Premeasures and algebras are respectively easier to handle than measures and σ-algebras. Carathéodory's theorem allows us to extend a premeasure μ_0 defined on an algebra $\mathcal{A}$ to a measure μ defined on $\mathcal{M}(\mathcal{A})$.

Theorem 2.17 (Carathéodory's Theorem). *Let $\mathcal{A}$ be an algebra on X and μ_0 a premeasure on $\mathcal{A}$. Then there exists a measure μ on $\mathcal{M}(\mathcal{A})$ such that $\mu|\mathcal{A} = \mu_0$. Furthermore, μ is unique on sets of finite measure.*

2.2.4. *Measurable Functions.* The development of measure theory was motivated by a desire to extend the theory of integration to a broader class of functions. The functions which turn out to be suitable are called measurable functions; they are an exact analogue of continuous functions for measure spaces.

Definition 2.18. Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces. A function $f : X \to Y$ is called (Σ_X, Σ_Y)-measurable (or just measurable if the σ-algebras are understood) if $f^{-1}(E) \in \Sigma_X$ for all $E \in \Sigma_Y$.

The following proposition is often useful in proving the measurability of a function.

Proposition 2.19. *Let $f : X \to Y$ be a function between measure spaces (X, Σ_X) and (Y, Σ_Y), and let $\Sigma_Y = \mathcal{M}(\mathcal{E})$. Then f is (Σ_X, Σ_Y)-measurable iff $f^{-1}(E) \in \Sigma_X$ for all $E \in \mathcal{E}$.*

Measurable functions give us one way of extending measures from one measurable space to another.

Definition 2.20. Let (X, Σ_X, μ_X) and (Y, Σ_Y, μ_Y) be measure spaces, and let $f : X \to Y$ be a measurable function. Then the set function

$$f_*\mu_X(E) = \mu_X(f^{-1}(E))$$

is a measure on Σ_Y, called the push-forward of μ_X by f. Likewise, the set function

$$f^*\mu_Y = \mu_Y(f(E))$$

is a measure on Σ_X, called the pull-back of μ_Y by f.

The names push-forward and pull-back are sometimes confusing. It helps to remember that they refer to the action of f on the measure, and not the set being measured.

2.3. Probability Theory.

2.3.1. *Basic Concepts.* Probability theory is a mathematical attempt to describe phenomena which involve chance or randomness. While it developed independently from measure theory, modern probability theory relies heavily on measure theoretic concepts. In particular, the set of all outcomes of a random process is associated with a space, and the events made up of subsets of these outcomes are associated with a σ-algebra on that space. A measure is then used to assign a probability of occurrence to each of these events.

Definition 2.21. A sample space or probability space is a measure space (Ω, B, P) with the property that $P(\Omega) = 1$. The measure P is called a probability measure, the elements of Ω are called outcomes, and the elements of B are called events. For $E \in B$, $P(E)$ is said to be the probability of the event E.

For any $E \in B$, $P(E)$ is a number between zero and one. We can think of it as the proportion of the time we can expect E to occur.

Often, it is useful to associate to each outcome in our sample space a numerical or other value. For example, we may like to associate a hand of blackjack with its total value; we could also assign a hand of blackjack a value of 1 if it contains a face card, and a value of 0 otherwise.

The set of values we assign to our outcomes is called the state space of our random process. It is also equipped with a σ-algebra. The map that takes outcomes to elements of our state space is called a random variable. The most common state space is $(\mathbb{R}, \mathcal{B}_{\mathbb{R}})$, so we define a random variable as real-valued.

Definition 2.22. A random variable X is a measurable function from a sample space (Ω, B, P) to $(\mathbb{R}, \mathcal{B}_{\mathbb{R}})$.

A real-valued random variable assigns a number to each outcome. The probability of an event can be calculated by adding up the values the random variable takes at each outcome in the event, in other words by integrating the random variable over the event. Every random variable also defines a probability measure on the real numbers.

Definition 2.23. Let (Ω, B, P) be a sample space, and let X be a random variable on Ω. For $E \in \mathcal{B}_{\mathbb{R}}$, let $P_X(E) = P(X^{-1}(E))$. Then P_X is a probability measure on $(\mathbb{R}, \mathcal{B}_{\mathbb{R}})$, called the distribution of X. The function $F_X(t) = P_X((-\infty, t])$ is called the distribution function of X.

Note that P_X is the push-forward of P by X. A family $\{X_\alpha\}_{\alpha \in A}$ is called identically distributed if $P_{X_\alpha} = P_{X_\beta}$ for all $\alpha, \beta \in A$.

A collection $X_1, ..., X_n$ of random variables on (Ω, B, P) defines a map X from (Ω, B, P) to $\mathbb{R}^n$ by

$$X(\omega) = (X_1(\omega), \ldots, X_n(\omega)).$$

This map, in turn, defines a probability measure on $\mathbb{R}^n$.

Definition 2.24. Let $X = (X_1, \ldots, X_n)$ be a collection of random variables on (Ω, B, P). The joint distribution of $X_1, \ldots, X_n$ is the unique probability measure on $\mathbb{R}^n$ which is defined as follows for a Borel rectangle $\prod_{i=1}^n B_i$, where $B_i \in \mathcal{B}_\mathbb{R}$ for each i:

$$P_{(X_1,\ldots,X_n)}\left(\prod_{i=1}^n B_i\right) = P\left(X^{-1}\left(\prod_{i=1}^n B_i\right)\right) = P\left(\bigcap_{i=1}^n X_i^{-1}(B_i)\right).$$

The concepts of mean and variance can be defined for any probability measure on $(\mathbb{R}, \mathcal{B}_\mathbb{R})$, and distributions allow us to define these concepts for arbitrary random variables as well.

Definition 2.25. Let P (P_X) be a probability measure on $(\mathbb{R}, \mathcal{B}_\mathbb{R})$ (distribution). Then the mean of P (P_X) is given by

$$\bar{P} = \int_\mathbb{R} t dP(t) \qquad \left(\bar{X} = \int_\mathbb{R} t dP_X(t)\right),$$

and the variance of P (P_X) is given by

$$\sigma_P^2 = \int_\mathbb{R} (t - \bar{P})^2 dP(t) \qquad \left(\sigma_X^2 = \int_\mathbb{R} (t - \bar{X})^2 dP_X(t)\right),$$

if these integrals exist.

2.3.2. *Independence.* One probabilistic concept that has no measure-theoretic analogue is that of independence.

Definition 2.26. Let (Ω, B, P) be a probability space, and let $E, F \in B$. The conditional probability of E with respect to F is

$$P(E|F) = \frac{P(E \cap F)}{P(F)}.$$

E and F are said to be independent if $P(E|F) = P(E)$.

The conditional probability $P(E|F)$ is the probability that E will occur if F does. E and F are independent if the occurence of F has no effect on how likely E is. We note that E and F are independent iff $P(E \cap F) = P(E)P(F)$. The concept of independence can also be extended to collections of events and random variables.

Definition 2.27. Let (Ω, B, P) be a probability space. The collection of events $\{E_\alpha\}_{\alpha \in A} \subset B$ is said to be independent if

$$P(E_{\alpha_1} \cap \ldots \cap E_{\alpha_n}) = \prod_{i=1}^n P(E_{\alpha_i})$$

for all $n \in \mathbb{N}$ and all distinct $(\alpha_1, \ldots, \alpha_n) \in A^n$.

Definition 2.28. Let $\{X_\alpha\}_{\alpha \in A}$ be a collection of random variables on the probability space (Ω, B, P). Then $\{X_\alpha\}_{\alpha \in A}$ is independent if for each $B \in \mathcal{B}_\mathbb{R}$ the set of events $\{X_\alpha^{-1}(B)\}_{\alpha \in A}$ are independent.

There is an equivalent characterization of independence for a collection of random variables that is often easier to work with.

Proposition 2.29. *Let $\{X_\alpha\}_{\alpha \in A}$ be a collection of random variables on the probability space (Ω, B, P). Then $\{X_\alpha\}_{\alpha\in A}$ is independent if for each finite subset $(\alpha_1, \dots, \alpha_n) \subset A$ and for each n, the joint distribution*

$$P_{(X_{\alpha_1},\dots,X_{\alpha_n})} = \prod_{i=1}^{n} P_{X_{\alpha_i}}.$$

In other words, the finite joint distributions of $\{X_\alpha\}_{\alpha\in A}$ factor.

2.3.3. *The Normal Distribution.* One of the most important distribution functions is the gaussian or normal distribution.

Definition 2.30. A random variable X is said to be normally distributed if it has distribution

$$P_X(E) = \int_E g(x; \bar{X}, \sigma_X)dx = \int_E \frac{1}{\sqrt{2\pi}\sigma_X} \exp\left[-\frac{(x-\bar{X})^2}{2\sigma_X^2}\right] dx,$$

and corresponding distribution function

$$F_X(t) = \int_{-\infty}^{t} g(x; \bar{X}, \sigma_X)dx.$$

$g(x; \bar{X}, \sigma_X)$ is called the gaussian with mean $\bar{X}$ and variance σ_X^2. The normal distribution with mean zero and variance one is called the standard normal distribution.

We will often abbreviate $g(x; 0, \sigma_X)$ by $g(x; \sigma_X)$. The importance of the normal distribution is in how it relates to other distributions. If a finite collection of n independent, identically distributed random variables is averaged together in a suitable way, the distribution of the average approaches the standard normal distribution as n grows large. This may explain why many natural random processes have a normal distribution; they are the result of many other random processes, occurring on smaller scales.

Theorem 2.31 (Central Limit Theorem). *Let $\{X_i\}_{i=1}^\infty$ be a family of independent, identically distributed random variables with mean $\bar{X}$ and variance σ_X^2. Define the random variable*

$$Z_n = \sum_{i=1}^{n} \frac{X_i - \bar{X}}{\sigma_X \sqrt{n}}.$$

Then for all t,

$$\lim_{n\to\infty} F_{Z_n}(t) = \lim_{n\to\infty} P(Z_n \le t) = \int_{-\infty}^{t} g(x; 0, 1)dx.$$

2.3.4. *Stochastic Processes.* A stochastic process is a mathematical model of a random process that evolves in time. It consists of a family of random variables indexed by time.

Definition 2.32. Let a probability space (Ω, Σ, P) and a set T be given. Then a stochastic process on $\Omega \times T$ is a family $\{X_t\}_{t\in T}$ of real-valued random variables on Ω.

The random variable X_t assigns a numerical value to each event in B at time t; its distribution P_{X_t} gives the probability that the system takes on certain real values at time t. The probability that the system takes on values in the Borel sets $B_1, \ldots, B_n$ at times $t_1, \ldots, t_n$ can be calculated as

$$P\left(\bigcap_{i=1}^{n} X_{t_i}^{-1}(B_i)\right).$$

This defines a unique probability measure on $\mathbb{R}^n$.

Definition 2.33. Let $\{X_t\}_{t\in T}$ be a stochastic process. The finite distributions of the process are the probability measures defined on Borel rectangles $\prod_{i=1}^{n} B_i$ by

$$P\left(\bigcap_{i=1}^{n} X_{t_i}^{-1}(B_i)\right),$$

for all $(t_1, \ldots, t_n) \in T^n$, and for all n.

A stochastic process is completely characterized by its finite distributions.

Given $t, s \in T$, the difference $X_t - X_s$ of random variables is itself a random variable, called an increment of the stochastic process.

Definition 2.34. A stochastic process $\{X_t\}_{t\in T}$ on (Ω, B, P) has independent increments if for all $s, t, u, v \in T$ with $s < t \le u < v$, the random variables $X_t - X_s$ and $X_v - X_u$ are independent.

If a stochastic process has independent increments, its finite distributions are easy to write down.

Proposition 2.35. *If a stochastic process $\{X_t\}_{t\in T}$ on (Ω, B, P) has independent increments, then its finite distributions are*

$$P_{(X_{t_1},\ldots,X_{t_n})} = \prod_{i=1}^{n} P_{X_{t_i}-X_{t_{i-1}}}.$$

3. The Wiener Process

The first step in quantifying the behavior of an average Brownian path is to derive a physically plausible but sufficiently abstract mathematical model of Brownian motion. We will model Brownian motion in one dimension as a stochastic process $\{X_t\}_{t\in[0,\infty)}$ on $(\mathbb{R}, \mathcal{B}_{\mathbb{R}}, dx)$, where dx denotes Lebesgue measure.

We assume for convenience that a Brownian particle always starts at the origin, and so we set $X_0 = 0$. Brownian motion results from a large number of random collisions, and the effects of each collision last only until the next collision (a very short time), so we want the increments of our stochastic process to be independent.

The collisions are as likely to occur on the left of the particle as on the right, and so the particle is as likely to move in one direction as the other. Thus we expect that on average the particle will not move at all, and we set the mean of X_t to be zero.

On the other hand, the more time elapses, the more collisions the particle undergoes, and the further it will be possible for the particle to move. Thus we would like the variance of X_t to be proportional to t.

Since Brownian motion is a natural process, we might expect its finite distributions to be normally distributed. A more formal argument is as follows. We can divide the interval $[s, t]$ into n equal subintervals $[t_0 = s, t_1], \ldots, [t_{n-1}, t_n = b]$. Then,

$$X_t - X_s = \sum_{j=1}^{n} (X_{t_j} - X_{t_{j-1}}).$$

Since our stochastic process has independent increments, this is a sum of independent, identically distributed random variables. The central limit theorem thus suggests that $X_t - X_s$ should be normally distributed.

Putting this all together, we obtain that for some C, $X_t - X_s$ should have the normal distribution given by $g(x; 0, C(t-s))$. C, called the diffusion constant, is related to the temperature of the solution, and describes the speed with which Brownian particles move. If $C = 1$, the stochastic process we have described is called an abstract Wiener process. Its finite distributions are defined on Borel rectangles $B = \prod_{i=1}^{n} B_i$ by

$$P_{(X_{t_1}, X_{t_2}, \ldots, X_{t_n})}(B) = \int_{B_1} \cdots \int_{B_n} \prod_{i=1}^{n} g(x; x_i, t_i - t_{i-1}) dx_n \ldots dx_1,$$

where $t_0 = x_0 = 0$.

4. Constructing Wiener Measure

4.1. **Setting the Stage.** We'll deal with Brownian motion in one dimension. This will make our analysis easier, and the extension to the multidimensional case is straightforward [2].

The space of Brownian paths in one dimension is a subset of the set of all functions from the positive real axis to the real numbers:

$$\Omega' = \{\omega : [0, \infty) \to \mathbb{R}\}.$$

The positive real axis can be thought of as time, and the value of the function or path at time t is the position of the Brownian particle at that time. The Brownian paths are the paths ω which are continuous, since a pollen grain cannot disappear in one location and reappear in another. We also specify that $\omega(0) = 0$. Dealing with the larger set Ω will simplify the construction, and, as we shall see in Section 5, we lose nothing by doing it.

As mentioned in the introduction, it will be convenient for us to modify this space in other ways as well. Firstly, we choose the range of our functions to be the one point compactification of the real numbers, $\dot{\mathbb{R}}$. Secondly, we take functions from $\mathbb{Q}^+ = \mathbb{Q} \cap [0, \infty)$ to $\dot{\mathbb{R}}$. This makes our function space a countable product of copies of $\dot{\mathbb{R}}$,

$$\Omega = \prod_{t \in \mathbb{Q}^+} \dot{\mathbb{R}}.$$

We equip this space with the product topology, τ. Recall that this is the topology generated by the sets

$$V = \pi_t^{-1}(U)$$

where $t \in \mathbb{Q}^+$, π_t is the projection of Ω onto the t^{th} copy of $\dot{\mathbb{R}}$, and U is an open subset of $\dot{\mathbb{R}}$. A Brownian path ω is in the set V if it passes through U at time t, that is if $\omega(t) \in U$. Note that, thanks to Tychonoff's theorem, (Ω, τ) is compact.

We want to define our measure on the Borel σ-algebra generated by the product topology, Σ_1. Our choice of function space is lucky; by restricting ourselves to rational times, we obtain that Σ_1 is equivalent to two other σ-algebras on Ω, one of which is easier to work with, and one of which is more powerful. The more manageable of these is the product over $\mathbb{Q}^+$ of the Borel σ-algebra on $\mathbb{R}$ (Definition 2.11),

$$\Sigma_2 = \bigotimes_{t \in \mathbb{Q}^+} \mathcal{B}_{\mathbb{R}}.$$

The more useful of these is the Borel σ-algebra generated by the topology of uniform convergence on Ω (Definition 2.5), which we denote by Σ_3. For general product spaces, these are not equivalent (see p.6).

Proposition 4.1. $\Sigma_1 = \Sigma_2$.

Proof. The product topology τ is generated by sets of the form $\pi_t^{-1}(U)$, where $t \in \mathbb{Q}^+$ and $U \in \tau_{\mathbb{R}}$. Since t is restricted to lie in $\mathbb{Q}^+$, τ can contain at most countable unions of sets of the above form. On the other hand, Σ_2 contains countable unions of sets of the form $\pi_t^{-1}(B)$, where $t \in \mathbb{Q}^+$ and $B \in \mathcal{B}_{\mathbb{R}}$. Since every open set is a Borel set, $\tau \subset \Sigma_2$. Thus, by Lemma 2.10, $\Sigma_1 \subset \Sigma_2$. To show the opposite inclusion, look at one of the generators of Σ_2. This set is contained in $\mathcal{M}(\pi_t^{-1}(\tau_{\mathbb{R}}))$, the σ-algebra generated by the inverse image under projection onto t of the usual topology on $\mathbb{R}$. Since $\pi_t^{-1}(\tau_{\mathbb{R}})$ is contained in the generating set of Σ_1, we can apply Lemma 2.10 again, completing the proof. □

Proposition 4.2. $\Sigma_1 = \Sigma_3$.

Proof. Let $B_\Omega(\omega, \epsilon)$ denote the ball of radius epsilon about the path $\omega \in \Omega$, in the metric of uniform convergence, and $B_{\mathbb{R}}(x, \epsilon)$ denote the ball of radius epsilon about $x \in \mathbb{R}$, in the Euclidean metric. Now,

$$B_\Omega(\omega, \epsilon) = \bigcup_{t \in \mathbb{Q}^+} \pi_t^{-1} B_{\mathbb{R}}(\omega(t), \epsilon).$$

But the latter set is a countable intersection of generators of Σ_1, and so an element of Σ_1. Thus by Lemma 2.10, Σ_1 contains Σ_3. For each $\omega \in \pi_t^{-1}(B_{\mathbb{R}}(x, \epsilon))$, define

$$d_\omega = \min\{|\omega(t) - (x - \epsilon)|, |\omega(t) - (x + \epsilon)|\},$$

the minimum distance between $\omega(t)$ and $B^c_{\mathbb{R}}(x, \epsilon)$. Then,

$$\pi_t^{-1}\left(B_{\mathbb{R}}(x, \epsilon)\right) = \bigcup_{\omega \in \pi_t^{-1}(B_{\mathbb{R}}(x,\epsilon))} B_\Omega(\omega, d_\omega).$$

The latter set is in the topology of uniform convergence, so by a second application of Lemma 2.10, we obtain our result. □

We will denote this (single) σ-algebra by Σ. To make the underlying product topology easier to read about, we will call sets that have a form similar to that of V above simple bundles. A simple bundle is a collection of paths (think of their graphs in the right half of $\mathbb{R}^2$) tied together by the fact that they all pass through a certain subset $E \subset \dot{\mathbb{R}}$ at a certain time. A bundle is a collection of paths tied together at n times by n subsets.

Definition 4.3. Let $t \in \mathbb{Q}^+$, $E \subset \dot{\mathbb{R}}$. Then the set of all paths passing through E at time t,

$$P(t, E) = \pi_t^{-1}(E)$$

is called a simple bundle. Let $T = (t_1, \ldots, t_n)$ be a vector of positive, rational numbers, and let $\mathcal{E} = (E_1, ..., E_n)$ be a collection of subsets E_i of $\dot{\mathbb{R}}$. Then $P(T, \mathcal{E})$, the collection of paths passing through E_i at time t_i for each i, is called a bundle.

Note that every simple bundle is a bundle. Also,

$$P(T, \mathcal{E}) = P(t_1, E_1) \cap \ldots \cap P(t_n, E_n).$$

If $E_i \in \mathcal{B}_{\mathbb{R}}$ for all i, we call $P(T, \mathcal{E})$ a Borel bundle. If each E_i is open, we call $P(T, \mathcal{E})$ an open bundle. The product topology is the topology generated by the open bundles. Two simple bundles $P(t, E)$, $P(t^{'}, E^{'})$ are disjoint iff $t = t^{'}$ and $E^{'} \cap E = \emptyset$, and two bundles $P(T, \mathcal{E})$, $P(T^{'}, \mathcal{E}^{'})$ are disjoint iff there is some t_i that is an element of both T and $T^{'}$, and $E_i^{'} \cap E_i = \emptyset$.

4.2. **Our Algebra and Premeasure.** To use Caratheodory's theorem, we must define a premeasure on an algebra which will generate Σ. We take our algebra, $\mathcal{A}$, to be the algebra generated by simple bundles in which E is a half-open interval of the form $(a, b]$ with $-\infty \leq a < b \leq \infty$. We will call these sets simple h-bundles; a finite intersection of simple h-bundles will be called an h-bundle. If $T = (t_1, \ldots, t_n) \in (\mathbb{Q}^+)^n$, $A = (a_1, \ldots, a_n) \in \dot{\mathbb{R}}^n$, and $B = (b_1, \ldots, b_n) \in \dot{\mathbb{R}}^n$, then we write

$$\bigcap_{i=1}^{n} P(t_i, (a_i, b_i]) = P(T, (A, B]).$$

Proposition 4.4. *Every $A \in \mathcal{A}$ is a finite union of disjoint h-bundles:*

$$A = \bigcup_{i=1}^{n} P(T_i, (A_i, B_i]).$$

Proof. All sets of this form are in our algebra, so the class of all such sets is contained in our algebra. We need to show that our algebra is contained in the class of these sets. All simple h-bundles have this form. We must show that the family of sets of this form is closed under complements and finite intersections; then, the algebra generated by the simple h-bundles will be contained in this family. As for finite intersections, let $A = \bigcup_{i=1}^{n} P_i$ and $B = \bigcup_{j=1}^{m} Q_j$ with P_i and Q_j h-bundles for each i, j. Then,

$$P \cap Q = \bigcup_{i,j} P_i \cap Q_j,$$

and $P_i \cap Q_j$ is an h-bundle. What's more, if either $i \neq k$ or $j \neq l$, then $P_i \cap Q_j$ is disjoint from $P_k \cap Q_l$. Thus, the union is disjoint. As for complements, the complement of a simple h-bundle is a union of simple h-bundles:

$$(P(t,(a,b]))^c = P(t,(-\infty,a]) \cup P(t,(b,\infty]) = P_-^c \cap P_+^c.$$

As for an h-bundle $P = \bigcap_{i=1}^n P_i$, we have

$$P^c = (P_1^c) \cup (P_1 \cap P_2^c) \cup \ldots \cup (P_1 \cap \ldots \cap P_n^c)$$

$$= \bigcup_{i=1}^{n} \left(P_{i+}^c \cap \bigcap_{j=1}^{i-1} P_j \right) \cup \bigcup_{i=1}^{n} \left(P_{i-}^c \cap \bigcap_{j=1}^{i-1} P_j \right) \tag{1}$$

This is a finite, disjoint union of $2n$ h-bundles. Now, let A have the proper form, with

$$A = \bigcup_{i=1}^{n} P(T_i,(A_i,B_i]) = \bigcup_{i=1}^{n} \bigcap_{j=1}^{n_i} P(t_j^i,(a_j^i,b_j^i]).$$

Let $T = \bigcup_{i=1}^n T_i$, and let $m = |T|$, the number of elements in T. Note that

$$P(T_i,(A_i,B_i]) = P(T_i,(A_i,B_i]) \cap P(T \setminus T_i,(-\infty,\infty]),$$

in other words we can rewrite the h-bundles making up A so that they all share the same vector of times T. So, we have

$$A = \bigcup_{i=1}^{n} P\left(T,\left(\widetilde{A}_i,\widetilde{B}_i\right]\right) = \bigcup_{i=1}^{n} \bigcap_{j=1}^{m} P(t_i,(a_j^i,b_j^i]).$$

Now, by (1), we can write

$$\left[P\left(T,\left(\widetilde{A}_i,\widetilde{B}_i\right]\right)\right]^c = \bigcup_{j=1}^{2m} Q_j^i,$$

where each Q_j^i is an h-bundle (not necessarily simple) and if $j \neq k$, $Q_j^i \cap Q_k^i = \emptyset$. Then, using de Morgan's laws, we have

$$A^c = \bigcap_{i=1}^{n} \bigcup_{j=1}^{2m} Q_j^i = \bigcup_{j=1}^{2m} \left[\bigcap_{i=1}^{n} Q_j^i\right].$$

The union on the right is disjoint because the Q_j^i are. By applying mathematical induction, we obtain that the class of sets with the specified form is closed under finite intersections. $\square$

Proposition 4.5. $\mathcal{M}(\mathcal{A}) = \Sigma$.

Proof. Observe that $P(t,(a,b]) = \pi_t^{-1}((a,b])$, and recall that $\mathcal{B}_{\mathbb{R}}$ is generated by the half-open intervals. Then, by Lemma 2.10, $\mathcal{M}(\mathcal{A}) = \bigotimes_{t\in\mathbb{Q}^+} \mathcal{B}_{\mathbb{R}}$. Finally, applying Proposition 4.1 gives us our result. $\square$

We define our premeasure on bundles composed of intervals of arbitrary topological status. (This will be useful later on.) In particular, it is defined on h-bundles and thus on $\mathcal{A}$. The premeasure accords with the finite distributions of the Wiener process as derived above.

Definition 4.6. Let $P = P(T, A, B) = \bigcap_1^n P(t_i, a_i, b_i)$. Then

$$\mu_0(P) = \int_{a_1}^{b_1} \cdots \int_{a_n}^{b_n} \prod_{i=1}^{n} g(x_i; t_i - t_{i-1})\, dx_n \ldots dx_1,$$

where $g(x; \sigma)$ is the gaussian distribution with variance σ^2, and $t_0 = 0$.

Note that μ_0 is finitely additive by definition on disjoint h-bundles.

4.3. **Extending μ_0 to Σ.** To use Carathéodory's theorem, we must show μ_0 is countably additive within $\mathcal{A}$. After reducing this to countable additivity on h-bundles, we show inequalities in both directions. The first, that the measure of a countable union of disjoint h-bundles is less than the sum of their measures, is an easy consequence of finite additivity. The second requires a more subtle argument. We use the "continuity" of our premeasure to approximate the summands and the sum by open and compact sets, respectively. A compactness argument then gives us our result.

Theorem 4.7. *There is a unique measure μ on Σ such that $\mu|\mathcal{A} = \mu_0$.*

Proof. Let $\{A_i\}_1^\infty \subset \mathcal{A}$, let $A = \bigcup_1^\infty A_i \in \mathcal{A}$, and let

$$A_i = \bigcup_{j=1}^{n_i} P(T_i^j, (A_i^j, B_i^j]), \qquad A = \bigcup_{k=1}^{n} P(T_k, (A_k, B_k]).$$

Write A_i as the finite disjoint union of h-bundles such that for each j, $P(T_i^j, (A_i^j, B_i^j]) \subset P(T_k, (A_k, B_k])$ for some k. Thus, the sequence $\{A_i\}_1^\infty$ can be partitioned into finitely many sequences of h-bundles, such that the union of the h-bundles in each sequence is one of the h-bundles composing A. Since μ_0 is finitely additive, we can consider each sequence of h-bundles separately; thus it suffices to consider a sequence of h-bundles whose union is an h-bundle. Let $\{P_i\}_1^\infty$ be such a sequence, and let $P = \bigcup_1^\infty P_i$. By finite additivity,

$$\begin{aligned}\sum_1^N \mu_0(P_i) &= \mu_0\left(\bigcup_1^N P_i\right) \\ &\leq \mu_0\left(\bigcup_1^N P_i\right) + \mu_0\left(\bigcup_N^\infty P_i\right) \\ &= \mu_0(P).\end{aligned}$$

Now, letting $N \to \infty$, we obtain that

$$\sum_1^\infty \mu_0(P_i) \leq \mu_0(P).$$

Let

$$P_i = P(T_i, (A_i, B_i]) = \bigcap_{j=1}^{n_i} P(t_j^i, (a_j^i, b_j^i]),$$

and let

$$P = P(T, (A, B]) = \bigcap_{k=1}^{n} P(t_k, (a_k, b_k]).$$

Let $\Delta = (\delta_1, \ldots, \delta_n)$, and let

$$\hat{P} = P(T, [A + \Delta, B - \Delta]) = \bigcap_{k=1}^{n} P(t_k, [a_k + \delta_k, b_k - \delta_k]).$$

Now, $\hat{P}$ is closed in the topology of pointwise convergence, since if $\{\omega_n\}_{n=1}^{\infty} \subset \hat{P}$ is such that ω_n converges to ω pointwise, then for $k = 1, \ldots, n$,

$$\omega(t_k) = \lim_{n\to\infty} \omega_n(t_k) \in [a_k + \delta_k, b_k - \delta_k].$$

As a closed subset of a compact space, $\hat{P}$ is compact. Also,

$$P \setminus \hat{P} \subset \bigcup_{k=1}^{n} P(t_k, (a_k, a_k + \delta_k]) \cup \bigcup_{k=1}^{n} P(t_k, (b_k - \delta_k, b_k]),$$

and so

$$\mu_0(P) - \mu_0(\hat{P}) = \sum_{k=1}^{n} \left(\int_{a_k}^{a_k+\delta_k} + \int_{b_k-\delta_k}^{b_k} \right) g(x; t_k) dx = \epsilon.$$

Since the integrand above is continuous (in fact, it is C^∞), it is clear that by taking the elements of Δ to be sufficiently small, we can make ϵ as small as we like. We denote the interior of an h-bundle P by $\overset{\circ}{P}$, and we note that

$$\overset{\circ}{P} = P(T, (A, B)) = \bigcap_{i=1}^{n} P(t_i, (a_i, b_i)),$$

with $\mu_0\left(\overset{\circ}{P}\right) = \mu_0(P)$. Now, $\left\{\overset{\circ}{P}_i\right\}_{i=1}^{\infty}$ is an open covering of $\hat{P}$. Thus there exists a finite subcovering $\left\{\overset{\circ}{P}_{n_i}\right\}_{i=1}^{N}$ of $\hat{P}$, so that

$$\hat{P} \subset \bigcup_{i=1}^{N} \overset{\circ}{P}_{n_i}.$$

Then,

$$\begin{aligned}
\mu_0(P) &= \mu_0\left(\hat{P}\right) + \epsilon \\
&\leq \sum_{i=1}^{m} \mu_0\left(\overset{\circ}{P}_{n_i}\right) + \epsilon \\
&\leq \sum_{i=1}^{\infty} \mu_0\left(\overset{\circ}{P}_i\right) + \epsilon \\
&= \sum_{i=1}^{\infty} \mu_0\left(P_i\right) + \epsilon
\end{aligned}$$

Letting $\epsilon \to 0$, we obtain the extension of μ_0 that we seek. To show that it is unique, note that the measure of Ω can be calculated as

$$\mu(P(t, \mathbb{R})) = \int_{\mathbb{R}} g(x; t) dx = 1,$$

where t is any element of $\mathbb{Q}^+$. So μ is finite, and thus unique. $\square$

We point out that (Ω, Σ, μ) is a probability space, as promised.

5. Brownian Paths are Hölder Continuous

5.1. **Wiener Measure Is Concentrated on $H_{\frac{1}{2}}$.** We have constructed Wiener measure on the space Ω of all functions from $\mathbb{Q}^+$ to $\dot{\mathbb{R}}$. Unfortunately, this space is, in some respects, too big. For example, we would like to exclude physically impossible discontinuous paths from consideration; indeed, this is the only way we will be able to extend functions in Ω to functions in $\mathbb{R}^{[0,\infty)}$. The content of this section is that we can ignore discontinuous functions with impunity. From a measure theoretic point of view, we lose nothing by considering only continuous paths. In fact, there is a stronger continuity condition that almost every Brownian path satisfies.

Definition 5.1. A real-valued function $f : X \to \mathbb{R}$ on a metric space X is called Hölder continuous of exponent $\alpha > 0$, or α-Hölder continuous, if there is a constant C such that for all $x, y \in X$

$$|f(x) - f(y)| < C\rho(x, y)^\alpha.$$

We denote the space of all such functions by H_α. We call $H = \bigcup_{\alpha>0} H_\alpha$ the space of Hölder continuous functions.

We remark first that if a path ω is β-Hölder continuous, it is also α-Hölder continuous for all $\alpha < \beta$ (we suppose that $\rho(x, y)$ is small). Thus, if $\beta > \alpha$, $H_\alpha \supset H_\beta$. In other words,

$$H_\alpha = \bigcup_{\beta \geq \alpha} H_\beta, \qquad H_\alpha^c = \bigcup_{\beta < \alpha} H_\beta.$$

As α grows to infinity, H_α shrinks to contain only constant functions. If we wanted to pick out the functions which are Hölder continuous with exponent α, but which violate Hölder continuity for all exponents $\beta > \alpha$, we could define

$$I_\alpha = H_\alpha \setminus \bigcup_{\beta > \alpha} H_\beta.$$

The theorem we will prove is that the probability of a Brownian path being Hölder continuous with exponenent $\alpha < 1/2$ is one. In other words, if $\alpha < 1/2$, then

$$\mu(H_\alpha) = 1.$$

What does this result mean? Imagine that alpha is decreasing from infinity, and the set H_α is expanding from the constant functions to include I_α for smaller and smaller α. So long as $\alpha > 1/2$, the measure of all these functions is zero. When we first include $I_{1/2}$, the measure of our set jumps to one, and so our set has full measure. As α decreases to zero, we include more and more functions, but these functions add nothing to the measure of our set, which continues to be one. Thus, μ experiences a δ-function like jump on $I_{1/2}$, and we say that it is concentrated on 1/2-Hölder continuous functions.

The intent of our statement is that almost all Brownian paths are in $I_{1/2}$, but the measure of $I_{1/2}$ is unkown. In fact, the most we can say is that for all $\epsilon > 0$,

$$\mu\left(\bigcup_{\alpha\in[\frac{1}{2}-\epsilon,\frac{1}{2}]} I_\alpha\right) = 1.$$

5.2. **Writing Down H_α.** To prove the Hölder continuity of Brownian paths, we will express H_α^c in terms of smaller sets whose Wiener measure is easy to estimate. This will be done by first writing H_α in terms of smaller sets. We will obtain an estimate for $\mu(H_\alpha^c)$ in terms of α, and we will show that this expression is zero only if $\alpha \in (0, 1/2)$.

A function that is continuous will be uniformly continuous on any finite interval $[0, s] \subset [0, \infty)$. It will have this property if for each $\epsilon > 0$, there is some $\delta > 0$ so that whenever $t, t' \in [0, s]$ and $|t - t'| \leq \delta$, $|\omega(t) - \omega(t')| \leq \epsilon$. With this as motivation, we let

$$F(s, \delta, \epsilon) = \bigcap_{\substack{t,t'\in[0,s]\\ |t-t'|\leq\delta}} \{\omega \in \Omega : |\omega(t) - \omega(t')| \leq \epsilon\}.$$

Hölder continuous functions have their growth within a small interval restricted uniformly by a multiple of a power of the size of that interval. Furthermore, this is true for intervals of all sizes. Thus,the set of functions that are α-Hölder continuous on $[0, s]$ with continuity constant c is given by the set

$$G(s, c) = \bigcap_{0<\delta\leq s} F(s, \delta, c\delta^\alpha).$$

But the continuity constant c can be anything, so the set of α-Hölder continuous functions on $[0, s]$ can be written as

$$H_\alpha(s) = \bigcup_{c>0} G(s, c),$$

and thus

$$H_\alpha = \bigcap_{s>0} H_\alpha(s) = \bigcap_{s>0}\bigcup_{c>0}\bigcap_{0<\delta\leq s} F(s, \delta, c\delta^\alpha).$$

Taking the complement and applying de Morgan's laws we find

$$H_\alpha^c = \bigcup_{s>0}\bigcap_{c>0}\bigcup_{0<\delta\leq s} F^c(s, \delta, c\delta^\alpha).$$

Proposition 5.2. *For all $\alpha \in (0, \infty)$, H_α is measurable.*

Proof. Note first that for arbitrary $t, t' \in \mathbb{Q}^+$ sets of the form

$$\{\omega \in \Omega : |\omega(t) - \omega(t')| \leq \epsilon\}$$

are closed in the topology of pointwise convergence, since if $\{\omega_n\}_{n=1}^\infty \subset \Omega$ is such that $\omega_n \to \omega$ pointwise, then

$$|\omega(t) - \omega(t')| = \lim_{n\to\infty} |\omega_n(t) - \omega_n(t')| \leq \epsilon.$$

But then $\bigcap_{0<\delta\leq s} F(s,\delta,\epsilon)$ is an intersection of closed sets, so it is closed itself. Next, if $s < s'$, then every path in $H_\alpha(s')$ is also in $H_\alpha(s)$. So, as s increases, the set $H_\alpha(s)$ decreases in size. Thus, we can write

$$H_\alpha = \bigcap_{s>0} H_\alpha(s) = \bigcap_{n=1}^{\infty} H_\alpha(n).$$

Similarly, if ω is α-Hölder continuous with constant c for some $c < c'$, then it is also α-Hölder continuous with constant c'. So the size of $G(n,c)$ increases with c, and we can write

$$H_\alpha(n) = \bigcup_{m=1}^{\infty} G(n,m).$$

Thus H_α is a countable intersection of countable unions of closed sets, and so it is measurable. □

5.3. **Estimating $\mu(F^c)$.** We must estimate the measure of the sets

$$F^c(s,\delta,\epsilon) = \bigcup_{\substack{t,t'\in[0,s] \\ |t-t'|\leq\delta}} \{\omega\in\Omega : |\omega(t)-\omega(t')| > \epsilon\}.$$

We will do this in three steps. Let $T = (t_1,\ldots,t_n)$ and $\epsilon > 0$ be given. Our first step is to estimate the measure of the set

$$A(T,\epsilon) = \bigcup_{t_i,t_j\in T} \{\omega\in\Omega : |\omega(t_j)-\omega(t_i)| > \epsilon\}.$$

Then, we will use this estimate to bound the measure of the set

$$A(a,b,\epsilon) = \bigcup_{t_i,t_j\in[a,b]} \{\omega\in\Omega : |\omega(t_j)-\omega(t_i)| > \epsilon\},$$

for given $0 < a < b < \infty$. This will be accomplished using continuity from below. Finally, we estimate the measure of $F^c(s,\delta,\epsilon)$ for certain values of s and δ using $A(a,b,\epsilon)$. Note that each of the above sets is open in the product topology, and so measurable.

Before beginning, we define the complementary error function,

$$\rho(\delta,\epsilon) = \int_{|x|>\epsilon} g(x;\delta)dx.$$

All our estimates will be in terms of this function. It is common knowledge that

$$\rho(\delta,\epsilon) \leq \frac{1}{\sqrt{2\pi\delta}}\left[\epsilon + \sqrt{\epsilon^2 + \frac{4}{\pi}\sqrt{2\delta}}\right]^{-1} \exp\left[-\frac{\epsilon^2}{2\delta}\right].$$

Also, we can use the complementary error function to exactly calculate the measure of certain sets:

$$\mu(\{\omega\in\Omega : |\omega(b)-\omega(a)| > \epsilon\}) = \rho(b-a,\epsilon).$$

Intuitively, as $\delta\to 0$, the mass of $g(x;\delta)$ becomes concentrated around zero. Thus, for fixed ϵ, $\rho(\delta,\epsilon)\to 0$ as $\delta\to 0$. Thus, we can expect Brownian paths to move very little over very small distances; this is the heart of the concept of continuity.

Lemma 5.3. *Let $T = (t_1, \ldots, t_n)$ with $0 < a = t_1 < \ldots < t_n = b < \infty$, and let $\epsilon > 0$. Then, using the same terminology as above,* $\mu(A(T, \epsilon)) \leq 2\rho\left(b - a, \frac{\epsilon}{4}\right)$.

Proof. If $|\omega(t_j) - \omega(t_i)| > \epsilon$, then

$$\begin{aligned}\epsilon &< |\omega(t_j) - \omega(t_i)| \\ &\leq |\omega(t_j) - \omega(t_1)| + |\omega(t_i) - \omega(t_1)|\end{aligned}$$

so either $|\omega(t_j)-\omega(t_1)| > \frac{\epsilon}{2}$ or $|\omega(t_i)-\omega(t_1)| > \frac{\epsilon}{2}$, and thus $A(T, \epsilon) \subset B(T, \epsilon)$ where

$$B(T, \epsilon) = \{\omega \in \Omega : \exists j \in \{1, \ldots, n\} \text{ with } |\omega(t_j) - \omega(t_1)| > \frac{\epsilon}{2}\}.$$

B is a union of open sets, and so it is measurable. Now, for all ω, either $|\omega(t_n) - \omega(t_1)| > \frac{\epsilon}{4}$, or $|\omega(t_n) - \omega(t_1)| \leq \frac{\epsilon}{4}$. If the latter is true and $\omega \in B(T, \epsilon)$, then

$$\begin{aligned}|\omega(t_n) - \omega(t_j)| \geq& |\omega(t_j) - \omega(t_1)| + |\omega(t_n) - \omega(t_1)| \\ >& \epsilon - |\omega(t_n) - \omega(t_1)| \\ \geq& \frac{\epsilon}{2} - \frac{\epsilon}{4} = \frac{\epsilon}{4}.\end{aligned}$$

Let

$$\begin{aligned}C(T, \epsilon) &= \{\omega \in \Omega : |\omega(t_n) - \omega(t_1)| > \frac{\epsilon}{2}\}, \\ D_j(T, \epsilon) &= \{\omega \in \Omega : |\omega(t_n) - \omega(t_j)| > \frac{\epsilon}{4}\}, \\ E_j(T, \epsilon) &= \{\omega \in \Omega : |\omega(t_j) - \omega(t_1)| > \epsilon \\ &\qquad \text{and if } i < j, \text{ then } |\omega(t_i) - \omega(t_1)| \leq \epsilon\}.\end{aligned}$$

Note that $C(T, \epsilon)$ and $D_j(T, \epsilon)$ are open sets in the product topology, and $E_j(T, \epsilon)$ is the intersection of an open set and a closed set, so all are measurable. Then by the above discussion,

$$B(T, \epsilon) \subset C(T, \epsilon) \cup \bigcap_{j=2}^{n} \left(D_j(T, \epsilon) \cap E_j(T, \epsilon)\right).$$

Now, $D_j(T, \epsilon)$ and $E_j(T, \epsilon)$ are independent events, since the position of a particle at time t_k does not depend on its position at times $t_1, \ldots, t_j$. Thus, we have

$$\mu(B(T, \epsilon)) \leq \mu(C(T, \epsilon)) + \sum_{j=2}^{n} \mu(D_j(T, \epsilon))\mu(E_j(T, \epsilon)).$$

But $\mu(C(T, \epsilon)) \leq \rho\left(b - a, \frac{\epsilon}{2}\right)$, and $\mu(D_j(T, \epsilon)) \leq \rho\left(b - a, \frac{\epsilon}{4}\right)$. Since ρ is a decreasing function of ϵ,

$$\mu(B(T, \epsilon)) \leq \rho\left(b - a, \frac{\epsilon}{4}\right)\left(1 + \sum_{2}^{n} \mu(E_j(T, \epsilon))\right) \leq 2\rho\left(b - a, \frac{\epsilon}{4}\right).$$

□

Note that the estimate above depends only on $b - a = t_n - t_1$, and not on the other elements of T.

Lemma 5.4. *Let $0 \leq a < b < \infty$. Then* $A(a, b, \epsilon) \leq 2\rho\left(b - a, \frac{\epsilon}{4}\right)$.

Proof. Let $\{q_n\}_{n\in\mathbb{N}}$ be an enumeration of $\mathbb{Q}\cap[a,b]$, and let $\mathbb{Q}_n = \{q_1, \ldots, q_n\}$. Then $A(\mathbb{Q}_n, \epsilon) \subset A(\mathbb{Q}_{n+1}, \epsilon)$ for all n, and

$$\bigcup_{n=1}^{\infty} A(\mathbb{Q}_n, \epsilon) = A(a, b, \epsilon).$$

Thus, by continuity from below, we have that

$$\mu(A(a, b, \epsilon)) = \lim_{n\to\infty} \mu(A(\mathbb{Q}_n, \epsilon)).$$

But from Lemma 5.3, we have $\mu(A(\mathbb{Q}_n, \epsilon)) \leq 2\rho\left(b-a, \frac{\epsilon}{4}\right)$, and so we obtain our result. □

Lemma 5.5. *Let $\epsilon > 0$ be given, and let $s, \delta > 0$ be such that $n = \frac{s}{\delta}$ is an integer. Then $\mu(F^c(s, \delta, \epsilon)) \leq 2n\rho\left(\delta, \frac{\epsilon}{4}\right)$.*

Proof. Let $\omega \in F^c(s, \delta, \epsilon)$, and let $t, t' \in [0, s]$ be such that $|t - t'| \leq \delta$ and $|\omega(t) - \omega(t')| > \epsilon$. Now, there is some $j \in \{1, \ldots, n\}$ so that both $|t - j\delta| \leq \delta$ and $|t' - j\delta| \leq \delta$. Furthermore,

$$\begin{aligned}\epsilon &< |\omega(t) - \omega(t')| \\ &\leq |\omega(t) - \omega(j\delta)| + |\omega(t') - \omega(j\delta)|\end{aligned}$$

so either $|\omega(t) - \omega(j\delta)| > \frac{\epsilon}{2}$ or $|\omega(t') - \omega(j\delta)| > \frac{\epsilon}{2}$. Thus, either $\omega \in A\left((j-1)\delta, j\delta, \frac{\epsilon}{2}\right)$ or $\omega \in A\left(j\delta, (j+1)\delta, \frac{\epsilon}{2}\right)$. So,

$$F^c(s, \epsilon, \delta) \subset \bigcup_{j=1}^{n} A\left((j-1)\delta, j\delta, \frac{\epsilon}{2}\right).$$

Taking the measure of both sides and applying Lemma 5.3, we obtain our estimate. □

5.4. **The Main Theorem.** Now we have the estimate that we need to prove our main result.

Theorem 5.6. *For any $\alpha < 1/2$, $\mu(H_\alpha) = 1$.*

Proof. Recall that

$$H_\alpha^c = \bigcup_{n=1}^{\infty} \bigcap_{c>0} \bigcup_{0<\delta\leq s} F^c(s, \delta, k\delta^\alpha),$$

and thus,

$$\mu\left(H_\alpha^c\right) \leq \sum_{n=1}^{\infty} \mu\left(H_\alpha^c(n)\right).$$

The size of $G^c(n, c)$ decreases with c, so by continuity from above,

$$\mu\left(H_\alpha^c(n)\right) = \lim_{c\to\infty} \mu\left(G^c(n, c)\right).$$

The size of $F^c(n, \delta, \epsilon)$ increases with δ (the longer a Brownian particle wanders, the further it tends to move) and decreases with ϵ (we don't expect Brownian particles to move arbitrarily large distances in a short time). Choose an integer m so that $\frac{n}{m+1} \leq \delta \leq \frac{n}{m}$. Then,

$$F^c\left(n, \delta, c\delta^\alpha\right) \subset F^c\left(n, \frac{n}{m}, c\left(\frac{n}{m+1}\right)^\alpha\right).$$

Thus, we can bound the set $G^c(n,c)$ from above by a countable union,

$$G^c(n,c) \subset \bigcup_{m=1}^{\infty} F^c\left(n, \frac{n}{m}, \frac{cn^\alpha}{(m+1)^\alpha}\right),$$

and we can write

$$\mu\left(G^c(n,c)\right) \le \sum_{m=1}^{\infty} \mu\left(F^c\left(n, \frac{n}{m}, \frac{cn^\alpha}{(m+1)^\alpha}\right)\right).$$

Applying Lemma 5.5 and our previously established bound on the complementary error function, we have that

$$\mu\left(F^c\left(n, \frac{n}{m}, \frac{cn^\alpha}{(m+1)^\alpha}\right)\right) \le 2n\rho\left(\frac{n}{m}, \frac{cn^\alpha}{4(m+1)^\alpha}\right)$$

$$\le 2(m+1)^\alpha \left[A + \sqrt{A^2 + \frac{8(m+1)^{2\alpha}}{mn}}\right]^{-1} \exp\left[-B\left(\frac{m}{(m+1)^{2\alpha}}\right)\right],$$

where

$$A = \frac{c\sqrt{\pi}}{4n^{1-\alpha}}, \qquad B = \frac{c^2 n^{2\alpha-1}}{32},$$

$$= 2(1+m^{-1})^\alpha c^{-1}\left[\widetilde{A} + \sqrt{\widetilde{A}^2 + \frac{8(1+m^{-1})}{c^2 n}}\right]^{-1} \exp\left[-\widetilde{B} m^{1-2\alpha}\right],$$

where

$$\widetilde{A} = \frac{\sqrt{\pi}}{4m^\alpha n^{1-\alpha}}, \qquad \widetilde{B} = \frac{c^2 n^{2\alpha-1}}{32(1+m^{-1})^{2\alpha}}.$$

Removing the terms containing a positive power of m in the denominator, we find that as m approaches infinity, the bound looks like

$$\sqrt{\frac{n}{2}} \exp\left[-\left(\frac{m^{1-2\alpha} c^2 n^{2\alpha-1}}{32}\right)\right].$$

If $\alpha < 1/2$, then as $m \to \infty$, the terms above go to zero fast enough to ensure convergence of the infinite series. If the infinite series converges, we can exchange the sum (in m) and the limit (in c) in the following bound on $\mu\left(H_\alpha^c(n)\right)$:

$$\lim_{c\to\infty} \sum_{m=1}^{\infty} 2(1+m^{-1})^\alpha c^{-1}\left[\widetilde{A} + \sqrt{\widetilde{A}^2 + \frac{8(1+m^{-1})}{c^2 n}}\right]^{-1} \exp\left[-\widetilde{B} m^{1-2\alpha}\right].$$

Each term in the sum clearly goes to zero as $c \to \infty$. Thus, $\mu(H_\alpha^c(n)) = 0$, and so $\mu(H_\alpha^c) = 0$. □

Corollary 5.7. *For all* $\epsilon > 0$, $\mu\left(\bigcup_{\alpha\in[1/2-\epsilon,1/2]} I_\alpha\right) = 1$.

Proof. For ease of notation, define

$$J_\epsilon = \bigcup_{\alpha\in[\frac{1}{2}-\epsilon,\frac{1}{2}]} I_\alpha = H_{(\frac{1}{2}-\epsilon)} \setminus \bigcup_{\alpha>\frac{1}{2}} H_\alpha.$$

Now,

$$\bigcup_{\alpha>\frac{1}{2}} H_\alpha = \bigcup_n H_{\left(\frac{1}{2}+\frac{1}{n}\right)},$$

and the latter set is measurable; thus J_ϵ is measurable. Then,

$$\begin{aligned}\mu(J_\epsilon) =& \mu\left(H_{\left(\frac{1}{2}-\epsilon\right)} \setminus \bigcup_{\alpha>\frac{1}{2}} H_\alpha\right)\\ =& \mu\left(H_{\left(\frac{1}{2}-\epsilon\right)}\right) - \mu\left(\bigcup_{\alpha>\frac{1}{2}} H_\alpha\right)\\ =& 1 - \lim_{\alpha\to\frac{1}{2}^+} \mu(H_\alpha) = 1.\end{aligned}$$

□

6. Extending Our Measure

One of the benefits of having our measure concentrated on continuous functions is that these functions are entirely determined by their values at rational times. This allows us to embed Ω in the larger space $\Omega^{'}$ in a natural way.

Definition 6.1. Recall that $\Omega = \prod_{t\in\mathbb{Q}^+} \dot{\mathbb{R}}$ and $\Omega^{'} = \prod_{t\in[0,\infty)} \dot{\mathbb{R}}$. Define $E : \Omega \to \Omega^{'}$ as follows.

(i) If $\omega \in \Omega$ is continuous,

$$E(\omega) = \begin{cases} \omega(t) & \text{if } t \in \mathbb{Q}^+, \\ \lim_{n\to\infty} \omega(t_n) & \text{if } t \in [0,\infty) \setminus \mathbb{Q}^+, \text{ where } \{t_n\}_1^\infty \subset \mathbb{Q}^+ \\ & \text{and } \lim_{n\to\infty} t_n = t. \end{cases}$$

(ii) If $\omega \in \Omega$ is discontinuous,

$$E(\omega) = 0,$$

where 0 is the path that is identically zero for all times t.

Now, using the embedding E and the concept of a push-forward (Definition 2.20), we can extend μ to the measurable space $(\Omega^{'}, \Sigma^{'})$, where $\Sigma^{'}$ is the σ-algebra generated by the topology of uniform convergence on $\Omega^{'}$. In order to do this, we need only verify that the map E is $(\Sigma, \Sigma^{'})$-measurable.

Proposition 6.2. *The embedding E, as given above, is $(\Sigma, \Sigma^{'})$-measurable.*

Proof. We will use Proposition 2.19, and the notation of Proposition 4.2. Recall that $B_\Omega(\omega, \epsilon)$ denotes the ball of radius epsilon about the path $\omega \in \Omega$ in the metric of uniform convergence. If $\omega \in \Omega^{'}$ is continuous, $E^{-1}(\omega)$ is just the restiction of ω to rational times; if ω is discontinuous, then $E^{-1}(\omega) = \emptyset$. So,

$$E^{-1}\left(B_{\Omega^{'}}(\omega, \epsilon)\right) = B_\Omega(\omega, \epsilon) \cap \Omega_c,$$

where Ω_c is the set of all continuous paths in Ω. Thus, we must show that $\Omega_c \in \Sigma$. Now, a real-valued function ω that is continuous on $[0, \infty)$

is continuous and bounded on $[0, n]$ for all n, so it is uniformly continuous on $[0, n]$ for all n. Uniform continuity on $[0, n]$ holds iff for every $j \in \mathbb{N}$, there is some $k \in \mathbb{N}$ so that $|\omega(t) - \omega(s)| \leq 1/j$ whenever $s, t \in [0, n]$ and $|t - s| \leq 1/k$. Thus,

$$\Omega_c = \bigcap_{n=1}^{\infty} \bigcap_{j=1}^{\infty} \bigcup_{k=1}^{\infty} F\left(n, k^{-1}, j^{-1}\right),$$

so $\Omega_c \in \Sigma$. □

Corollary 6.3. *The set function*

$$\mu^{'}(S) = E_{*}\mu(S) = \mu(E^{-1}(S))$$

is a measure on $(\Omega^{'}, \Sigma^{'})$.

Since $\mathbb{Q}^+$ is dense in $[0, \infty)$, if $\omega \in \Omega^{'}$ is α-Hölder continuous, $E^{-1}(\omega) \in H_\alpha$ also. Thus, if $\alpha < 1/2$,

$$\mu^{'}(H_\alpha) = \mu(E^{-1}(H_\alpha)) = \mu(H_\alpha) = 1,$$

so $\mu^{'}$ is also concentrated on $I_{1/2}$.

7. Conclusions

We have constructed a version of Wiener measure using elementary (though not trivial) measure theory, in particular relying on Carathéodory's theorem. We have also shown that this version of Wiener measure is concentrated on Hölder continuous paths with exponent $\alpha < 1/2$, an important property of classical Wiener measure.

Taking the first step of constructing μ on the set of functions from $\mathbb{Q}^+$ to $\mathbb{R}$ turned out to be essential. The fact that our functions are defined only at rational times is used in Proposition 4.1, Proposition 4.2, and Lemma 5.4, and so also in the results relying on them, including Proposition 4.5 and Theorem 5.6. The "continuity" of our premeasure was the only thing necessary to apply Carathéodory's theorem. To be more precise, our premeasure μ_0 was regular on $\mathcal{A}$, where regularity is the property that for all $A \in \mathcal{A}$,

$$\begin{aligned}\mu_0(A) =& \sup\{\mu_0(K) : K \text{ is compact and } K \subset A\} \\ =& \sup\{\mu_0(U) : U \text{ is open and } A \subset U\}.\end{aligned}$$

I suspect that a compactness argument similar to the one we used could be applied to any premeasure defined on a subset of a Borel σ-algebra and regular on its algebra.

A number of open questions about our measure remain. One is whether or not it is regular on Σ. Every probability measure constructed with the Riesz representation theorem is automatically regular; no such property is guaranteed for measures constructed using Carathéodory's theorem. Questions like those introduced at the beginning of the paper, about the average behavior of Brownian paths, have been answered for classical Wiener measure. Analogous results for our measure $\mu^{'}$ have not been presented here, and I leave the calculation of these probabilities as an exercise for beginning students of measure theory.

I would like to thank Daniel Ueltschi for proposing this paper, and helping to shape it. My thanks also go to Doug Pickrell for his advice.

References

[1] Chowdhury, D. 100 years of Einstein's theory of Brownian motion: from pollen grains to protein trains. arXiv:cond-mat/0504610v1, April 24, 2005.

[2] Folland, G.B. *Real Analysis: Modern Techniques and Their Applications* (2nd ed.). Wiley Interscience, New York (1999).

[3] Ginibre, J. Some Applications of Functional Integration in Statistical Mechanics, pp. 327-427, *Statistical Mechanics and Quantum Field Theory.* C. DeWitt & R. Sora eds., Gordon and Breach Science Publishers, New York (1971).

[4] Mazo, R.M. *Brownian Motion: Fluctuations, Dynamics and Applications.* Oxford University Press, New York (2002).

Wiener Measure Using Carathéodory's Theorem